arXiv:1404.0703v6 [cs.DB] 22 Dec 2015

Joins via Geometric Resolutions: Worst-case and Beyond

Mahmoud Abo Khamis[*‡] Hung Q. Ngo[*‡] Christopher Ré[†] Atri Rudra[*]

[*] Department of Computer Science and Engineering
University at Buffalo, SUNY
{mabokham,hungngo,atri}@buffalo.edu

[†] Department of Computer Science
Stanford University
chrismre@cs.stanford.edu

[‡] LogicBlox Inc.
{mahmoud.abokhamis,hung.ngo}@logicblox.com

Abstract

We present a simple geometric framework for the relational join. Using this framework, we design an algorithm that achieves the fractional hypertree-width bound, which generalizes classical and recent worst-case algorithmic results on computing joins. In addition, we use our framework and the same algorithm to show a series of what are colloquially known as beyond worst-case results. The framework allows us to prove results for data stored in Btrees, multidimensional data structures, and even multiple indices per table. A key idea in our framework is formalizing the inference one does with an index as a type of geometric resolution; transforming the algorithmic problem of computing joins to a geometric problem. Our notion of geometric resolution can be viewed as a geometric analog of logical resolution. In addition to the geometry and logic connections, our algorithm can also be thought of as backtracking search with memoization.

1 Introduction

Efficient processing of the natural join operation is a key problem in database management systems [1, 40, 70]. A large number of algorithms and heuristics for computing joins have been proposed and implemented in database systems, including Block-Nested loop join, Hash-Join, Grace, Sort-merge, index-nested, double pipelined, PRISM, etc. [12, 14, 32, 38]. In addition to their role in database management, joins (or variants) are powerful enough to capture many fundamental problems in logic and constraint satisfaction [39, 57], or subgraph listing problems [50, 51] which are central in social [68, 69] and biological network analysis [47, 60].

Not surprisingly, there has been a great deal of work on joins in various settings. A celebrated result is Yannakakis' algorithm, which shows that acyclic join queries can be computed in linear time [72] in data complexity (modulo a log factor). Over the years, this result was generalized to successively larger classes of queries based on various notions of widths: from *treewidth* (tw) [22, 62], *degree of acyclicity* [36, 37], *query width* (qw) [15], to *generalized hypertree width* (ghw) [31, 64]. From the bound of Atserias, Grohe and Marx [6, 35] (AGM bound), and its algorithmic proof [50], we recently know that there is a class of join algorithms that are optimal in the worst case, in the sense that for each join query the algorithm runs in time linear in the size of the worst-case output [50, 51, 71]. Combining a worst-case optimal join algorithm with Yannakakis' algorithm yields an algorithm running in time $\tilde{O}(N^{\text{fhtw}} + Z)$, where fhtw stands for *fractional hypertree width* [35], a more general notion than the widths mentioned above, and Z is the output size.

However, worst-case can be pathological. For example, input relations are typically already pre-processed and stored in sophisticated indices to facilitate fast query answering (in even *sub-linear* time). Motivated by this, recent work has gone beyond worst-case analysis to notions that are closer to *instance or pointwise* optimality. These beyond worst-case results have as their starting point the work of Demaine et al. [25] and Barbay and Kenyon [7, 8], who designed beyond worst-case algorithms for set intersection and union problems, which were recently extended to join processing [49].[1]

As one might expect, the algorithms that achieve the above varied results are themselves varied; they make a wide range of seemingly incompatible assumptions: data are indexed or not; the measures are worst-case or instance-based; they may rely on (or ignore) detailed structural information about the query or cardinality information about the underlying tables. With all this variety, our first result may be surprising: *we recover all of the above mentioned results with a single, simple algorithm.*[2]

Our central algorithmic idea is to cast the problem of evaluating a join over data in indices as a geometric problem; specifically, we reduce the join problem to a problem (defined below) in which one covers a rectangular region of a multidimensional space (with dimension equal to the number of attributes of the join) with a set of rectangular boxes. These boxes represent regions in the space in which we know output tuples are *not* present. Such rectangles are a succinct way to represent the information conveyed by these data structures. We illustrate these ideas by an example.

Example 1.1. Consider the relation $R(A, B) = \{4\} \times \{1, \ldots, 7\} \cup \{1, \ldots, 7\} \times \{4\}$, which is illustrated in Figure 1a. For now assume that R is stored in a B-tree with attribute order (A, B). Any two consecutive tuples (a, b_1) and (a, b_2) in R give rise to the box with its A side being the singleton interval $[a, a]$ and its other side being the interval (b_1, b_2). For example the gap box between the tuples $(4, 4)$ and $(4, 5)$ in Figure 1a is denoted by the blue box in Figure 1b. Consecutive tuples (a_1, b_1) and (a_2, b_2) for $a_1 \neq a_2$ give rise to bigger rectangles. Figure 1b illustrates all the gap boxes generated from this R. Suppose we want to compute the join $R(A, B) \bowtie S(B, C)$ for some other relation S. Then, the gap boxes from R will span all values in the C-dimension. Similarly, the gap boxes for S span all values in the A-dimension. And the output tuples are precisely the tuples (a, b, c) which do not fall into any gap boxes, from both indices of R and S.

Throughout this paper we will think of the data as integers for convenience, but our results assume only that the domains of attributes are discrete and ordered. For technical reasons, we will assume that the boxes are *dyadic boxes*, i.e., rectangles whose endpoints and side lengths can be encoded as powers of 2. Importantly, a dyadic interval can be

[1] This algorithm has been implemented in a commercial database system, LogicBlox, with promising but initial results. In our preliminary experimental results, the new algorithm on some queries on real social network data showed up to two, even three orders of magnitude speedup over several existing commercial database engines [52].

[2] We are unable to recover the more recent notions of widths, in particular, the notion of *submodular width* [45]. See Section 2 for a more detailed discussion.

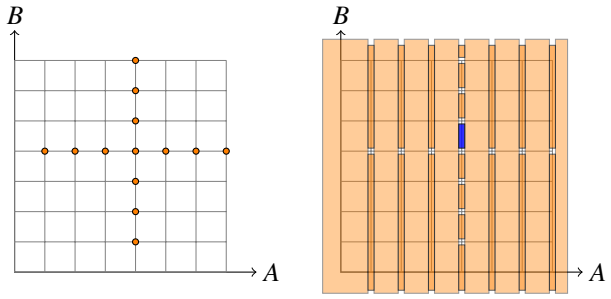

(a) Tuples from $R(A, B)$

(b) Gap boxes from (A, B)-ordered Btree

Figure 1: A relation and the corresponding gap boxes from sorted order (A, B).

thought of as a bitstring, which allows many geometric operations such as containment and intersection to be reduced to string operations that take time linear in the length of strings and so logarithmic in the size of the data. This encoding does increase the number of gap boxes, but by only a polylogarithmic factor in the input data size. With this idea, the central problem in this work is the *box cover problem*, which informally is defined as follows. (Formal definition is in Section 3.)

> Given a set of dyadic boxes $\mathcal{A}$, i.e. the gaps from the data,[3] our goal is to list all the points that are not covered by any box in $\mathcal{A}$.

The core of our algorithm solves essentially the boolean version of the box cover problem, where in addition to the set of boxes $\mathcal{A}$, one is also given a target box $\mathbf{b}$ and the goal is to check if $\mathbf{b}$ is covered by the union of boxes in $\mathcal{A}$. Our algorithm for the boolean box cover problem is recursive, with the following steps: We first check if any box $\mathbf{a} \in \mathcal{A}$ contains $\mathbf{b}$. If such $\mathbf{a}$ exists, we return it as a witness that $\mathbf{b}$ is covered. If not, then we split the box $\mathbf{b}$ into two halves $\mathbf{b}_1$ and $\mathbf{b}_2$ and recurse. In the recursive steps, we either find a point in the target box that is not covered (in which case we return it as a witness that the target box is not covered), or we discover two boxes $\mathbf{w}_1, \mathbf{w}_2 \in \mathcal{A}$ that contain $\mathbf{b}_1$ and $\mathbf{b}_2$ respectively. We then construct a single box $\mathbf{w}$ by combining $\mathbf{w}_1$ and $\mathbf{w}_2$ such that $\mathbf{w}$ contains $\mathbf{b}$, we add $\mathbf{w}$ to $\mathcal{A}$ [4], and we return it as a witness that $\mathbf{b}$ is covered. This algorithm needs to answer three questions:

- *How to find a box* $\mathbf{a} \in \mathcal{A}$ *containing the target box* $\mathbf{b}$ *if such a box exists?* This search procedure should be efficient, ideally in polylogarithmic time in the data size. Dyadic encoding of gap boxes makes this goal possible. Our algorithm stores boxes in $\mathcal{A}$ in a (multilevel) dyadic tree data structure.
- *How to split the input box* $\mathbf{b}$ *into* $\mathbf{b}_1$ *and* $\mathbf{b}_2$*?* A first natural scheme is to go in a fixed attribute order. We show that this scheme is sufficient to recover all the results mentioned earlier in this section. However, we also show that this approach is fundamentally limited. In particular, we show a novel alternate scheme that is able to achieve much stronger per-instance guarantees.
- *How to combine* $\mathbf{w}_1$ *and* $\mathbf{w}_2$ *to form* $\mathbf{w}$*?* The combine operation has two competing goals: it should be *complete* in that it can infer $\mathbf{b}$ (or a box that contains $\mathbf{b}$) and it should be *efficient* in that it should take at most polylogarithmic time in the data. For that purpose, we introduce a notion called *geometric resolution* (Figure 6). We show that this framework is complete and the resolution operation can be implemented efficiently as a simple operation on bitstrings. In conjunction with the efficient search procedure, this implies that the running time of the algorithm is the number of such resolutions (up to polylogarithmic factors in the data size). Thus, we can reason about the geometry of these covers instead of the algorithmic steps.[5]

In the first contribution, we show that our algorithm – named Tetris– is able to recover the worst-case algorithmic results shown in the top-half of Table 1, and the recent beyond-worst-case results of [49].

[3] We note that our algorithms assume an oracle access to $\mathcal{A}$ and some of our algorithms essentially minimize the number of accesses to the

Join Query Type	Run time	Recovers(R)/Subsumes(S)
WORST-CASE RESULTS		
α-acyclic	$N + Z$	Yannakakis [72] (R)
Arbitrary	$N + \mathsf{AGM}$	[50, 71] (R)
bounded 'width' w*	$N^{w*} + Z$	[15, 22, 31, 50, 71, 72] (S)
CERTIFICATE BASED RESULTS		
treewidth w	$\lvert C_\square\rvert^{w+1} + Z$	New and [49] (S)
treewidth 1	$\lvert C_\square\rvert + Z$	New

Table 1: Overview of our upper bounds achieved by the *same* algorithm called Tetris. The run times are up to poly-logarithmic factors and either in terms of the total input size N or in the size of the optimal certificate $C_\square$ as well as the output size Z. In the above, the bounded width $w^* = \text{fhtw} \leqslant \text{ghw} \leqslant \text{qw} \leqslant \text{tw} + 1$. Our worst-case result on fractional hypertree width (fhw) implies the other worst-case results on various notions of widths. Our result for treewidth w queries subsumes that of [49] since the latter only works for indices with mutually consistent sort orders. AGM denotes the AGM-bound for the query [6].

Our second contribution is to use these insights to go beyond known results. In previous work on beyond-worst-case analysis, one made an assumption that indexes were consistent with a single global ordering of attributes; a constraint that is not often met in practice. Our first results remove this restriction, which we believe argues for the power of the above framework. In particular, we reason about multiple Btrees on the same relation, multidimensional index structures like KD-trees and RTrees, and even sophisticated dyadic trees. In turn, this allows us to extend beyond-worst case analysis to a larger set of indexing schemes and, conceptually, this brings us closer to a theory of how indexing and join processing impact one another.

The idea of beyond worst-case complexity is captured by a natural notion of geometric certificate. In particular, a minimum-sized subset $C_\square \subseteq \mathcal{A}$ whose union is the same as the union of all input gap boxes in $\mathcal{A}$ is called a *gap box certificate* for the join problem. For beyond worst-case results, $|C_\square|$ is the analogous quantity to input size N that is used in the the worst-case results.

There are several reasons for our current certificate framework to use only 'gap' boxes and not input tuples (or more generally boxes that contain the input tuples). First, gap boxes directly generalize the results from [49], where it was shown that $|C_\square|$ is in the same order as the minimum number of comparisons that a comparison-based join algorithm has to perform in order to be certain that the output is correct. Second, we expect the input data to be very sparse in the ambient space. In particular, we show in this paper that $|C_\square| = O(N)$ and there are classes of input instances for which $|C_\square| = o(N)$ (or even $O(1)$). Third, gap boxes in some sense capture differences between different input indices. The same relation indexed in different ways give different sets of gap boxes which can all be used in evaluating the join. Last but not least, our move to use gaps rather than the input tuples themselves has a strong parallel with using proof by contradiction to prove logical statements. In hindsight, this parallel is precisely what results in the strong connection between our framework and resolution (indeed resolution is a specific form of proof by contradiction).

We show that for queries with treewidth 1 (i.e. query graphs are forests), we can compute them in $\tilde{O}(|C_\square| + Z)$ time[6], where Z is the output size. For general treewidth w join queries, we obtain a weaker runtime of $\tilde{O}(|C_\square|^{w+1} + Z)$.

We also develop a new and intriguing result, where we obtain a runtime of $\tilde{O}(|C_\square|^{n/2} + Z)$ for a query with n attributes. This subsumes and greatly extends the results on 3-cliques from previous work to all queries on n attributes (including n-cliques). Moreover, it shows that in some sense the clique queries remain the hardest for beyond worst-case analysis as well. Our geometric framework plays a crucial role in this result, both in the analysis and the design of our algorithm.

---

oracle.

[4] $\mathcal{A}$ is assumed to be a global variable: all levels of the recursion access the same $\mathcal{A}$. Check Section 4.2.1 for more details.

[5] This should be contrasted with traditional logical resolution that can potentially require $\Omega(N)$ time for a single step due to large clauses. At a high level, logical resolution is resolving *combinatorial* rectangles while our notion resolves geometric rectangles. See Appendix J for more discussion and examples.

[6] In this paper $\tilde{O}$ will hide poly-log N factors as well as factors that just depend on the query size, which is assumed to be a constant.

Finally, we also use our framework to provide lower bounds on the number of geometric resolutions that any algorithm needs. In particular, we consider three variants of geometric resolution in this paper. The most general kind (which is not as powerful as general logical resolution) resolves geometric boxes, which we call Geometric Resolution. We are able to recover all of the results in Table 1 with a weaker form of geometric resolution called Ordered Geometric Resolution, which corresponds to geometric resolution but when we only combine boxes in a fixed attribute order. We also consider an even more special case Tree Ordered Geometric Resolution, which corresponds to ordered geometric resolution when we do *not* cache the outcome of any resolution. Figure 2 summarizes where our upper and lower bounds fit in these classes of resolution.

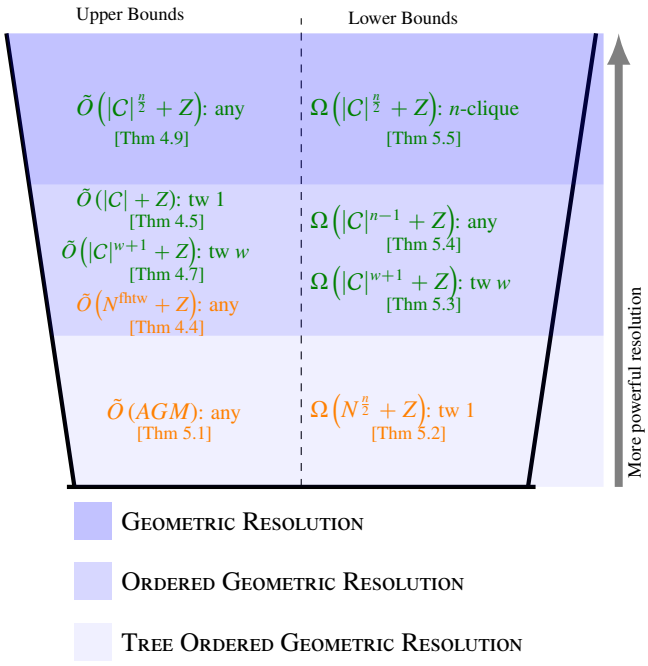

Figure 2: An overview of our results and the resolution framework. Bounds for worst-case complexity are denoted by orange (where N is the size of the largest relation) and certificate-based results are presented in green (where C is the optimal certificate). AGM denotes the bound on the output due to AGM and Z denotes the size of the output (per-instance basis). The bounds are presented in format time:query, where *any* denotes an arbitrary query on n attributes, *tw* w denotes a query on n attributes with treewidth w ($1 < w < n-1$) and *n-clique* denotes the n-variable clique query.

There is an intriguing connection between our framework and DPLL with clause learning used for #SAT. We address this further in Section 4.2.4.

2 Related Work

In a seminal work [72], Yannakakis showed that if the query is acyclic (or more precisely α-acyclic in Fagin's terminology [27]) then it can be evaluated in time $\tilde{O}(N+Z)$, where N is the input size (in terms of data complexity), and Z is the output size. Researchers have expanded the classes of tractable queries using an increasingly finer structural measure called the 'width' of the query, measuring how 'far' from being acyclic a query is. If the query 'width' is bounded by a constant, then the problem is tractable in the data complexity sense. The width notion progressed from *treewidth* (tw) [22, 62], *degree of acyclicity* [36, 37], *query width* (qw) [15], *hypertree width* and *generalized hypertree width* [31, 64].

Worst-case optimal join algorithms Atserias, Grohe, and Marx (AGM henceforth) [6, 35] derived a bound on the output size (the number of tuples in the output) using *both* the structural information about the query *and* the input relation sizes. Their bound (see Appendix A) is a function of the input relation sizes and a fractional edge cover of the hypergraph representing the query. By solving a linear program, we can obtain the best possible bound for the output size. We refer to this best bound as the *AGM bound*. AGM also showed that their bound is tight (in data complexity) by constructing a family of instances for which the output size is in the order of the bound. Similar but slightly weaker bounds were proven by Alon [4] and Friedgut and Kahn [29]. All of these results were proved via entropy-based arguments.

An algorithm whose worst-case runtime matches the AGM bound would be *worst-case optimal*. Such an algorithm was derived by Ngo, Porat, Ré, and Rudra (NPRR henceforth). Soon after, the Leapfrog Triejoin algorithm [71] was shown to run within the AGM bound. An even simpler but generic skeleton of a class of join algorithms which generalized both NPRR and Leapfrog Triejoin was shown to run within AGM bound [51].

Marx introduced yet another more general notion of width called *adaptive width* [44], which is equivalent to *submodular width* [43], and we are unable to recover this tighter notion of width using the results of this paper.

Beyond worst-case for joins with Minesweeper Beyond worst-case analysis in databases was formalized by Fagin et al.'s algorithm [28] for searching scored items in a database. These per instance guarantees are desirable, though they are very hard to achieve: there have been relatively few such results [2]. More relevantly, for the sorted set intersection problem, Demaine, López-Ortiz, and Munro [25], with followups by Barbay and Kenyon [7, 8] devised the notion of a *certificate* or a *proof*, which is a set of comparisons necessary to certify that the output is correct. An algorithm running in time proportional to the minimum certificate size (up to a log-factor and in data complexity) can be considered instance-optimal among comparison-based algorithms.[7]

The work of Demaine et al. and Barbay et al. was extended to general join queries [49] by defining the notion of a *comparison certificate* for a join problem, which roughly speaking is a set of propositional comparison statements about the input, such that two inputs satisfy the same set of propositional statements if and only if they have the same output. Intuitively, the minimum size of a comparison certificate is the minimum amount of work a comparison-based join algorithm has to do to correctly compute the output. A major technical assumption needed in prior work [49] was that all relations are indexed by BTrees according to a single *global attribute order* (GAO) index. For example, if the GAO is A, B, C, D (attributes participating in the query), and $R(A, C)$ is an input relation, then the BTree/trie for R has to branch on A before C. In this work, we are able to handle more general indexes (KD-trees, dyadic trees and multiple indices per relation) and do not require this assumption. To the best of our knowledge, the current work and that in [49] are the only two instances that present (near) instance optimal results for a large class of problems.

The analysis from [49] implies that we can use a 'box certificate' in place of a 'comparison certificate' because a box certificate has size at most the size of a comparison certificate (see Appendix B). This result inspired our investigation into the world of geometric certificates in this paper. Indeed we were able to generalize the results from [49] because the box certificates we considered in this paper are more general than the GAO-consistent boxes in [49].

Connections to DPLL As we will see in Section 4.2.4, Tetris is essentially a version of the DPLL algorithm. We would like to stress that the novelty of our work is to (i) adapt this well-known framework to a geometric view of joins and (ii) prove sharp bounds on the run time of Tetris.

Klee's measure problem in computational geometry A variant of the box cover problem (Definition 3.4) is the *Boolean* box cover problem (Definition 3.5): given a set $\mathcal{B}$ of n-dimensional boxes, determine whether their union covers the entire space. The Boolean box cover problem is a special case of *Klee's measure problem*: given a set $\mathcal{B}$ of n-dimensional boxes, compute the measure of their union. Klee's measure problem was solved by Overmars and Yap [56] in time $O(|\mathcal{B}|^{n/2}\log(|\mathcal{B}|))$, and later by Chan [13] in time $O(|\mathcal{B}|^{n/2})$. One corollary of this paper (Corollary F.12) shows that Klee's measure problem over the Boolean semiring can be solved in time $\tilde{O}(|\mathcal{C}_\square|^{n/2})$, where $\mathcal{C}_\square$ is any box certificate for $\mathcal{B}$. (By Definition 3.4, $|\mathcal{C}_\square| \leqslant |\mathcal{B}|$ and there are instances where $|\mathcal{C}_\square|$ is unboundedly smaller than

[7] As was observed in [49, 63], the log factor loss is necessary when dealing with comparison-based algorithms.

$|\mathcal{B}|$.) We also present tighter upper bounds for the box cover problem (and hence for Klee's measure problem over the Boolean semiring) in multiple special cases that are common in database joins (e.g. bounded tree-width, acyclicity, GAO-consistency...). Moreover, while the upper bound of $O(|\mathcal{B}|^{n/2})$ has not been shown to be tight for Klee's measure problem for $n \geqslant 3$, we show that our upper bounds (including $\tilde{O}(|\mathcal{C}_\square|^{n/2})$) are tight for all algorithms that are based on Geometric Resolution (see Corollary G.11).

3 Preliminaries

We give an overview of the strong connection between indices and gap boxes in Section 3.1 and then move on to our geometric notion of certificates in Section 3.2. We formally define our main geometric problem BCP in Section 3.3.

3.1 Gap boxes and indices

We informally describe the idea of gap boxes that capture database indices. The set of gap boxes depends intimately on the indices that store the relations. For example, for the relation in Figure 1a, Figure 1b shows the gap boxes generated by a BTree that uses the sort order (A, B). Figure 3a shows the gap boxes for the same relation when stored in a BTree

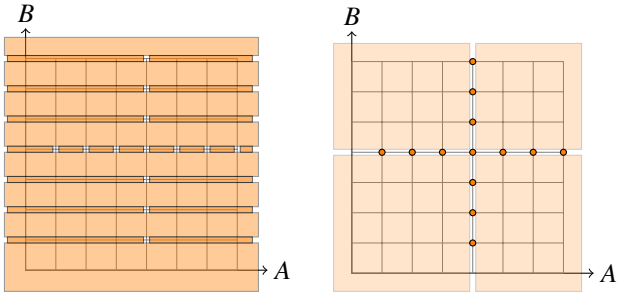

(a) (B, A)-consistent gap boxes (b) Gap-boxes in R from quad-trees

Figure 3: The gap boxes for the relation in Figure 1a from sorted order (B, A) and from a quad-tree type index.

with sort order (B, A). Note that the different sort order manifests itself in a completely different set of boxes. Finally, Figure 3b represents the boxes for the same relation when stored in a quad-tree. In addition to a completely different set of boxes from those in Figures 1b and 3a, the number of boxes is also much smaller. We will see another example for three attributes soon.

3.2 Geometric Certificates

For any attribute A, $\mathbf{D}(A)$ denotes its domain. For any join query Q, let atoms(Q) denote the set of constituting relations; in other words, we can write Q as $Q = \bowtie_{R \in \text{atoms}(Q)} R$. For any relation R (Q included), vars(R) denotes the set of its attributes.

We assume each input relation R is already indexed using some data structure that satisfies the following property. The data structure stores a collection $\mathcal{B}(R)$ of gap boxes whose union contains all points in $\prod_{A \in \text{vars}(R)} \mathbf{D}(A)$ which are *not* tuples in R. Note that there can be multiple indices per relation. Gap boxes from all those indices contribute to $\mathcal{B}(R)$. By filling out the coordinates not in vars(R) with "wild cards" (i.e. each one of those coordinates spans the entire dimension), we can without loss of generality view $\mathcal{B}(R)$ as a collection of gap boxes in the output space $\prod_{A \in \text{vars}(Q)} \mathbf{D}(A)$.

We begin with the notion of certificate. The size of the smallest such certificate will replace the input size as the measure of complexity of an instance in our beyond worst-case results.

Definition 3.1 (Box certificate)**.** A *box certificate* for Q is a subset of gap boxes from $\bigcup_{R \in \text{atoms}(Q)} \mathcal{B}(R)$ that cover every tuple not in the output. We use $\mathcal{C}_\square$ to denote a box certificate of minimum size for the instance.

We would like to stress the point above that the size of the smallest box certificate is intimately tied to the kind of index being used. In particular, for certain instances the certificate sizes might be much smaller for more powerful kinds of indices. This should be contrasted with the worst-case results of [50, 71] where BTrees with a single sort order are enough to obtain the optimal worst-case results (and using more powerful indices like quad-trees does not improve the results). Further, our algorithms do not assume the knowledge of $C_\square$ though they implicitly compute a box certificate C such that $|C| = \tilde{O}(C_\square)$.

For a more thorough discussion of indices, gap boxes, various notions of certificates and how they relate to box certificates, please see Appendix B. In particular, we can show that the notion of a box certificate is finer than the notion of comparison-based certificate used in [49].

Dyadic boxes For simplicity, but without any loss of generality, let us assume the domain of each attribute is the set of all binary strings of length d, i.e. $\mathbf{D}(A) = \{0,1\}^d$, for every $A \in \text{vars}(Q)$. This is equivalent to saying that the domain of each attribute is the set of all integers from 0 to $2^d - 1$. Since d is the number of bits needed to encode a data value of the input, d is logarithmic in the input size.

Definition 3.2 (Dyadic interval)**.** A *dyadic interval* is a binary string x of length $|x| \leqslant d$. This interval represents all the binary strings y such that $|y| = d$ and x is a prefix of y. Translating to the integral domain, the dyadic interval represents all integers in the range $[x2^{d-|x|}, (x+1)2^{d-|x|} - 1]$. Here, we view x as the integer corresponding to the bit string. The empty string $x = \lambda$ is a dyadic interval consisting of all possible values in the domain. (This serves as a wild-card.) If $|x| = d$, then it is called a *unit dyadic interval*, which represents a point in the domain.

Definition 3.3 (Dyadic box)**.** Let $\text{vars}(Q) = \{A_1, \ldots, A_n\}$. A *dyadic box* is an n-tuple of dyadic intervals: $\mathbf{b} = \langle x_1, \ldots, x_n \rangle$. If all components of $\mathbf{b}$ are unit dyadic intervals, then $\mathbf{b}$ represents a *point* in the output space. The dyadic box is the set of all tuples $\mathbf{t} = (t_1, \ldots, t_n) \in \prod_{i=1}^n \mathbf{D}(A_i)$ such that t_i belongs to the dyadic interval x_i, for all $i \in [n]$.

Note again that some dyadic intervals can be λ, matching arbitrary domain values; also, a dyadic box $\mathbf{b}$ contains a dyadic box $\mathbf{b}'$ if each of $\mathbf{b}$'s components is a prefix of the corresponding component in $\mathbf{b}'$. The set of all dyadic boxes forms a partially ordered set (poset) under this containment.

It is straightforward to show that every (not necessarily dyadic) box in n dimensions can be decomposed into a disjoint union of at most $(2d)^n = \tilde{O}(1)$ dyadic boxes. In particular, for every box certificate, there is a dyadic box certificate of size at most a factor of $\tilde{O}(1)$ larger. Henceforth, we will assume that all boxes are dyadic boxes. This assumption is also crucial for the discovery of an optimal box certificate. In particular, the number of dyadic boxes containing a given tuple is always at most $\tilde{O}(1)$. (See Appendix B.3 for the details.)

3.3 The Box Cover Problem

We assume the input index data structure(s) for an input relation R can return in $\tilde{O}(1)$-time the set of all dyadic gap boxes in $\mathcal{B}(R)$ containing a given tuple in $\prod_{A \in \text{vars}(R)} \mathbf{D}(A)$. This assumption holds for most of the common indices in relational database management systems such as BTree or trie. The objective of a general join algorithm is to list the set of all output tuples. Our join algorithm will attempt to take full advantage of the gaps stored in the input indices: it tries to compute/infer a collection of dyadic boxes whose union contains all tuples in $\prod_{A \in \text{vars}(Q)} \mathbf{D}(A)$ except the output tuples. (The smallest such collection is called a (dyadic) box certificate as defined in Definition 3.1.) Recall that an output tuple *is also* a (unit) dyadic box. Hence, the output dyadic boxes and the gap boxes together fill the entire output space. Consequently, we can think of a join algorithm as an algorithm that tries as fast as possible to fill up the entire output space with dyadic boxes of various shapes and sizes.

Abstracting away from the above idea, we first define a problem called the *box-cover problem* (or BCP).

Definition 3.4 (Box Cover Problem)**.** Given a set $\mathcal{A}$ of (dyadic) boxes, list all tuples *not* covered by any box in $\mathcal{A}$, i.e. list all tuples $\mathbf{t}$ such that $\mathbf{t} \notin \mathbf{b}$ for every $\mathbf{b} \in \mathcal{A}$. Define the (box) *certificate* for the instance $\mathcal{A}$ of BCP to be the smallest subset $C_\square \subseteq \mathcal{A}$ such that $\bigcup_{\mathbf{b} \in C_\square} \mathbf{b} = \bigcup_{\mathbf{b} \in \mathcal{A}} \mathbf{b}$.

Definition 3.5 (Boolean Box Cover Problem)**.** Given a set $\mathcal{A}$ of (dyadic) boxes, determine whether their union covers the entire output space, i.e. $\bigcup_{\mathbf{b} \in \mathcal{A}} \mathbf{b} = \langle \lambda, \ldots, \lambda \rangle$.

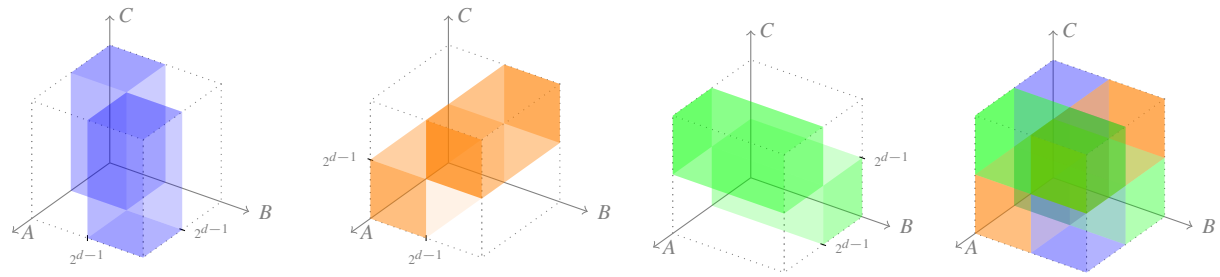

(a) Gap boxes for R (b) Gap boxes for S (c) Gap boxes for T (d) Union of all gap boxes

Figure 4: The first three figures show the gap boxes for $R(A,B), S(B,C)$ and $T(A,C)$ where each relation has the tuples in the dyadic boxes $0 \times 1 \cup 1 \times 0$, where recall $b \in \{0,1\}$ denotes the dyadic interval corresponding to all length d bitstrings that start with b. The last figure shows the union of all the gap boxes. Since the boxes cover all of the bounding box (which is denoted by the dashed gray box), the output is empty.

Given a join query Q, $\mathcal{B}(Q)$ denotes the set of all gap boxes from the input indices, i.e. $\mathcal{B}(Q) = \bigcup_{R \in \text{atoms}(Q)} \mathcal{B}(R)$. The following is straightforwardly true.

Proposition 3.6. *On input $\mathcal{A} = \mathcal{B}(Q)$, the output of* BCP *is exactly the same as the output of the join query Q. And, $|C_\square(\mathcal{B}(Q))| = |C_\square(Q)|$.*

We illustrate the connection between the triangle query $Q_\Delta = R(A,B) \bowtie S(B,C) \bowtie T(A,C)$ and the corresponding instance for the BCP. Consider the instance for Q_Δ in which R has pairs $(a,b) \in \{0,1\}^d \times \{0,1\}^d$ such that the first bits (or MSBs) of a and b are complements of each other. In this case, the gaps in R can be represented in a dyadic tree as depicted in Figure 4a: there are two gap boxes corresponding to all triples (a,b,c) such that the first bits of a and b are 0 and 1 respectively.[8] Further, let $(b,c) \in S$ ($(a,c) \in T$ resp.) if and only if the first bits of b and c (a and c resp.) are different. The corresponding gap boxes are depicted in Figures 4b and 4c. Then the BCP instance corresponds to the six gap boxes the union of which covers the entire output space (as depicted in Figure 4d), since the output of Q_Δ is empty for the given instance. See Figure 5 for another instance for the same join query when the output is non-empty.

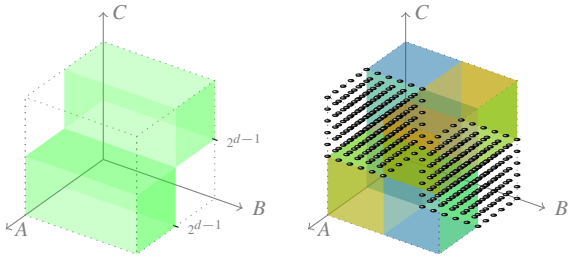

(a) Gap boxes for T' (b) Output points marked

Figure 5: Boxes for BCP instance corresponding to $R \bowtie S \bowtie T'$ with R and S as in Figure 4 and $T' = 0 \times 0 \cup 1 \times 1$ (boxes for T' shown on the left). The union of the boxes along with the output tuples are specifically marked out in the right figure.

Definition 3.7 (Support of a dyadic box)**.** Let $\mathbf{b} = \langle x_1, \ldots, x_n \rangle$ be a dyadic box. Its *support*, denoted by $\mathsf{support}(\mathbf{b})$, is the set of all attributes A_i for which $x_i \neq \lambda$. It follows that, if $\mathbf{b} \in \mathcal{B}(R)$ for some relation R, then the $\mathsf{support}(\mathbf{b}) \subseteq \text{vars}(R)$.

[8] By contrast one gap box in Figure 4a would correspond to roughly 2^{d-1} gap boxes if R was stored in a BTree.

Definition 3.8 (Supporting hypergraph of a set of boxes)**.** Let $\mathcal{A}$ be a collection of dyadic boxes. The *supporting hypergraph of* $\mathcal{A}$, denoted by $\mathcal{H}(\mathcal{A})$, is the hypergraph whose vertex set is the set $\mathcal{V}$ of all attributes participating in boxes of $\mathcal{A}$, and whose edge set is the set of all $\mathsf{support}(\mathbf{b})$, $\mathbf{b} \in \mathcal{A}$.

Proposition 3.9. *Let Q be any join query, and* tw *denote* tree-width*, then* $\mathsf{tw}(\mathcal{H}(\mathcal{B}(Q))) \leqslant \mathsf{tw}(Q)$.

Proof. For every box $\mathbf{b} \in \mathcal{B}(Q)$, we have $\mathsf{support}(\mathbf{b}) \subseteq \mathrm{vars}(R)$ for some $R \in \mathrm{atoms}(Q)$. Thus, every edge of the hypergraph $\mathcal{H}(\mathcal{B}(Q))$ is a subset of some edge of the hypergraph of Q. This means every tree decomposition of the hypergraph of Q is a tree decomposition of the hypergraph $\mathcal{H}(\mathcal{B}(Q))$. The proposition follows trivially. □

A dyadic segment x is *non-trivial* if $x \neq \lambda$ and x is not a unit segment. Let $\mathcal{A}$ be a set of dyadic boxes on attribute set $\mathcal{V} = \{A_1, \ldots, A_n\}$. A *global attribute order* (GAO) is an ordering σ of attributes in $\mathcal{V}$.

Definition 3.10 (GAO-consistent boxes)**.** Let $\mathcal{A}$ be a set of dyadic boxes on $\mathcal{V}$ and σ be a GAO on $\mathcal{V}$. Then, $\mathcal{A}$ is said to be σ*-consistent* if the following conditions are met: (a) For every box $\mathbf{b} = \langle x_1, \ldots, x_n \rangle \in \mathcal{A}$, there is at most one x_i for which x_i is non-trivial. (b) For every box $\mathbf{b} = \langle x_1, \ldots, x_n \rangle \in \mathcal{A}$, if x_i is non-trivial, then $x_j = \lambda$ for all j such that A_j comes after A_i in σ.

Note that if Q is a join query whose input relations are indexed consistently with a GAO σ, then $\mathcal{B}(Q)$ is σ-consistent. (See [49] for the definition of GAO-consistent indices. In short, the search tree for each relation is indexed using an attribute order consistent with the GAO.)

4 Upper Bounds

We formally define the notion of geometric resolution in Section 4.1. Our main algorithm Tetris is presented in Section 4.2. We present rederivations of existing results using Tetris in Section 4.3 (worst-case results) and in Section 4.4 (beyond-worst case results, which recover and generalize results from [49]). Finally, we present our new beyond worst-case result that works for arbitrary queries in Section 4.5.

4.1 Geometric Resolution

Our algorithm uses the framework of *geometric resolution*, which is a special case of logical resolution. The two input *clauses* to geometric resolution are two dyadic boxes, say,

$$\mathbf{w}_1 = \langle y_1, \ldots, y_n \rangle \text{ and } \mathbf{w}_2 = \langle z_1, \ldots, z_n \rangle$$

that have to satisfy the following two properties: (1) There exists a position $\ell \in [n]$ and a string x such that $y_\ell = x0$ and $z_\ell = x1$ (where x can be λ and xb denotes the concatenation of string x and bit b); and (2) For every other $j \in [n] \backslash \{\ell\}$, either y_j is a prefix of z_j or z_j is a prefix of y_j.

The result of the geometric resolution or the *resolvent* is the dyadic box

$$\mathbf{w} = \langle y_1 \cap z_1, \ldots, y_{\ell-1} \cap z_{\ell-1}, x, y_{\ell+1} \cap z_{\ell+1}, \ldots, y_n \cap z_n \rangle,$$

where we use $y_i \cap z_i$ to denote the *longer* of the two strings y_i, z_i. Geometrically, $\mathbf{w}_1$ and $\mathbf{w}_2$ are adjacent in the ℓth dimension, and in the other dimensions we are taking the intersection of those two dyadic segments which are contained in one another. For the rest of the paper, unless we explicitly mention otherwise, whenever we say resolution we mean geometric resolution. Pictorially this can be visualized for $n = 2$ as in Figure 6.

We briefly explain the name 'resolution.' In propositional logic, the resolution of two clauses C_1 and C_2 is a clause C_3 such that every truth assignment satisfying both C_1 and C_2 must satisfy C_3. The geometric resolution of two boxes $\mathbf{w}_1$ and $\mathbf{w}_2$ is a box $\mathbf{w}_3$ such that every point *not* covered by either $\mathbf{w}_1$ nor $\mathbf{w}_2$ must not be covered by $\mathbf{w}_3$.

Proposition 4.1 (Completeness of geometric resolution)**.** *Given a set of boxes $\mathcal{A}$ such that the union of all the boxes in $\mathcal{A}$ covers some box* $\mathbf{b}$*, there exists a sequence of geometric resolutions on $\mathcal{A}$ that results in a box* $\mathbf{b}'$ *that contains* $\mathbf{b}$.

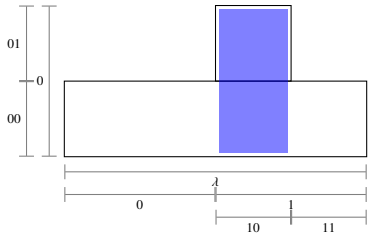

Figure 6: Geometric resolution on the vertical axis between two dyadic rectangles $\langle \lambda, 00 \rangle$ (bottom box) and $\langle 10, 01 \rangle$ (top box). The resolution result ($\langle 10, 0 \rangle$) is highlighted and is drawn slightly smaller than its correct size for illustration purposes.

The crux of this paper is to show that one can efficiently find a small sequence of geometric resolutions that solves BCP. The inputs to most of the resolutions made by our algorithms will have an even more restricted structure:

$$\mathbf{w}_1 = \langle y_1 \ , \ \ldots \ , \ y_{\ell-1} \ , \ x_\ell 0 \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle \tag{1}$$

$$\mathbf{w}_2 = \langle z_1 \ , \ \ldots \ , \ z_{\ell-1} \ , \ x_\ell 1 \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle, \tag{2}$$

where for every $i < \ell$, either y_i or z_i is a prefix of the other.

Definition 4.2 (Ordered geometric resolution)**.** Given two dyadic boxes $\mathbf{w}_1$ and $\mathbf{w}_2$ of the format shown in (1) and (2), the *ordered geometric resolution* of $\mathbf{w}_1$ and $\mathbf{w}_2$ is the dyadic box

$$\mathbf{w} = \langle y_1 \cap z_1 \ , \ \ldots \ , \ y_{\ell-1} \cap z_{\ell-1} \ , \ x_\ell \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle \tag{3}$$

We say that $\mathbf{w}$ is the result of resolving $\mathbf{w}_1$ and $\mathbf{w}_2$ *on attribute* A_ℓ. (Note that x_ℓ might be λ.)

4.2 The Algorithm

Our algorithm for BCP at its core solves essentially the boolean version of the BCP using a sub-routine called TetrisSkeleton. The sub-routine is then repeatedly invoked by the *outer* algorithm – Tetris – to compute the output of the BCP instance.

4.2.1 The Core Algorithm

The Boolean version of BCP is the following problem: given a set of dyadic boxes $\mathcal{A}$ and a target box $\mathbf{b}$, determine if $\mathbf{b}$ is covered by the (union of) boxes in $\mathcal{A}$. TetrisSkeleton solves this problem by not only answering YES or NO, but also generating an *evidence* for its answer:

- If $\mathbf{b}$ is covered by $\mathcal{A}$, then output a box $\mathbf{w}$ that covers $\mathbf{b}$ such that $\mathbf{w}$ is covered by the union of boxes in $\mathcal{A}$.
- If $\mathbf{b}$ is not covered by $\mathcal{A}$, then output a point/tuple in $\mathbf{b}$ that is not covered by any box in $\mathcal{A}$.

TetrisSkeleton has a very natural recursive structure. We fix a splitting attribute order (SAO) of the query, say $(A_1, \ldots, A_n)$. Following this order, we find the first dimension on which $\mathbf{b}$ is *thick* and thus can be split into two halves $\mathbf{b}_1$ and $\mathbf{b}_2$. If we can find an uncovered point in either half, then we can immediately return. Otherwise, we have recursively found two boxes $\mathbf{w}_1$ and $\mathbf{w}_2$, each of which covers one half of $\mathbf{b}$. Each box may not cover $\mathbf{b}$ as a whole. Hence, we resolve the two boxes $\mathbf{w}_1$ and $\mathbf{w}_2$ by creating a maximal box $\mathbf{w} \subseteq \mathbf{w}_1 \cup \mathbf{w}_2$, making sure that $\mathbf{w}$ covers both $\mathbf{b}_1$ and $\mathbf{b}_2$; hence, $\mathbf{w}$ covers $\mathbf{b}$. Figure 7 illustrates the main idea.

TetrisSkeleton is presented in Algorithm 1. There are three extra things that Algorithm 1 does over the basic outline above. First, we handle the base cases when $\mathbf{b}$ is already covered by a box in $\mathcal{A}$ itself in lines 1 and 2 (see Appendix C.1 on how we can implement this step in $\tilde{O}(1)$ time using a *multi-level dyadic tree* data structure) and when

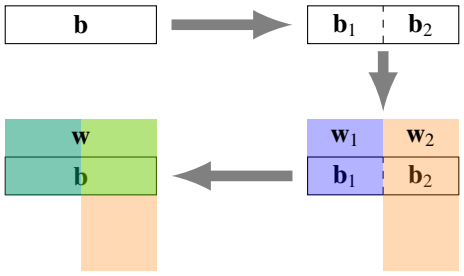

Figure 7: Illustration of main resolution step.

b is a unit box that is not covered by any box in $\mathcal{A}$ (and hence cannot be covered by any boxes derived from $\mathcal{A}$ either) in lines 3 and 4. Second, we check boundary conditions in lines 11 and 16. Finally, in line 19, we add back the result of resolution from line 18 to $\mathcal{A}$. (The last step is crucial in proving most of our results.) We defer a more detailed discussion on Resolve and Split-First-Thick-Dimension to the end of this section.

4.2.2 The Outer Algorithm: Tetris

The TetrisSkeleton algorithm was designed for the boolean BCP. We now present the simple idea that allows us to use TetrisSkeleton as a sub-routine and solve the general BCP problem. The input to the general BCP problem is a set of boxes $\mathcal{B}$ to which we have oracle access. The *oracle* represents the pre-built database indices of input relations from a join query. In particular, given a unit box $\mathbf{w}$, the oracle can return the set of boxes in $\mathcal{B}$ containing $\mathbf{w}$ in $\tilde{O}(1)$-time. (See Appendix C.1 for more.)

Algorithm 2, named Tetris, solves BCP by continuously calling TetrisSkeleton on input $\mathcal{A}$, called the knowledge base, with the target box being the universal box $\mathbf{b} = \langle \lambda, \ldots, \lambda \rangle$. We will explain how different initializations of $\mathcal{A}$ lead to different guarantees in later sections.

After each invocation of TetrisSkeleton, $\mathcal{A}$ is amended with a few more boxes and the next invocation of TetrisSkeleton is on the enlarged knowledge base $\mathcal{A}$. Apart from resolvents that are cached by TetrisSkeleton in $\mathcal{A}$, Tetris amends $\mathcal{A}$ with two types of boxes: output (unit) boxes and boxes from $\mathcal{B}$. To be more specific, if TetrisSkeleton returns $(\textsc{true}, \mathbf{w})$, then we know there are no tuples to output and we can stop. However, if TetrisSkeleton returns $(\textsc{false}, \mathbf{w})$, then we check if $\mathbf{w}$ is *not* covered by any box in $\mathcal{B}$. If so, we know $\mathbf{b}$ is an output point and we can output that point and amend $\mathcal{A}$ with $\mathbf{b}$. Otherwise we know that $\mathcal{A}$ was not properly initialized in which case we amend $\mathcal{A}$ with boxes in $\mathcal{B}$ that cover $\mathbf{b}$ and repeat.

4.2.3 Recursion and Resolution

We next flesh out the two key operations that were not specified in the description of Algorithm 1: how to split a box $\mathbf{b}$ into two halves in line 6, and how to resolve two *witnesses* $\mathbf{w}_1$ and $\mathbf{w}_2$ in line 18 of TetrisSkeleton (Algorithm 1).

We first explain what the Split-First-Thick-Dimension routine does. Consider a dyadic box $\mathbf{b} = \langle x_1, x_2, \ldots, x_n \rangle$. If $|x_i| = d$, then x_i represents a unit dyadic segment, which corresponds to a flat slice through the A_i-dimension. The box $\mathbf{b}$ is flat and is not splittable along such dimension. The first thick dimension is the smallest value $\ell \in [n]$ for which $|x_\ell| < d$. Because $\mathbf{b}$ is not a unit box, there must exist such an ℓ. In that case, the call

$$(\mathbf{b}_1, \mathbf{b}_2) \quad \leftarrow \quad \text{Split-First-Thick-Dimension}(\mathbf{b})$$

in line 6 of Algorithm 1 returns the following pair:

$$\begin{aligned} \mathbf{b}_1 &= \langle x_1 \ , \ \ldots \ , \ x_{\ell-1} \ , \ x_\ell 0 \ , \ x_{\ell+1} \ , \ \ldots \ , \ x_n \rangle \\ \mathbf{b}_2 &= \langle x_1 \ , \ \ldots \ , \ x_{\ell-1} \ , \ x_\ell 1 \ , \ x_{\ell+1} \ , \ \ldots \ , \ x_n \rangle. \end{aligned}$$

Note again that by definition $|x_i| = d$ for all $i < \ell$. It is easy to implement the above bitstring operation in $O(dn)$ time, which by our convention is $\tilde{O}(1)$ time.

Algorithm 1 TetrisSkeleton ($\mathbf{b}$)

Precondition: A global set of boxes $\mathcal{A}$ ▷ Our knowledge base
Precondition: Pick a splitting attribute order (SAO) $(A_1, \ldots, A_n)$
Input: Target box $\mathbf{b}$
Output: A pair $(v, \mathbf{w})$, where $\mathbf{w}$ is a cover box for $\mathbf{b}$ if v is TRUE, and an uncovered point if v is FALSE

1: **If** there is a box $\mathbf{a} \in \mathcal{A}$ such that $\mathbf{a} \supseteq \mathbf{b}$ **then**
2: **Return** $(\text{TRUE}, \mathbf{a})$
3: **else If** $\mathbf{b}$ is a unit box **then**
4: **Return** $(\text{FALSE}, \mathbf{b})$
5: **else**
6: $(\mathbf{b}_1, \mathbf{b}_2) \leftarrow$ Split-First-Thick-Dimension($\mathbf{b}$)
7: ▷ Cut $\mathbf{b}$ into two equal halves
8: $(v_1, \mathbf{w}_1) \leftarrow$ TetrisSkeleton($\mathbf{b}_1$)
9: **If** v_1 is FALSE **then**
10: **Return** $(\text{FALSE}, \mathbf{w}_1)$
11: **else If** $\mathbf{w}_1 \supseteq \mathbf{b}$ **then**
12: **Return** $(\text{TRUE}, \mathbf{w}_1)$
13: $(v_2, \mathbf{w}_2) \leftarrow$ TetrisSkeleton($\mathbf{b}_2$) ▷ Backtracking
14: **If** v_2 is FALSE **then**
15: **Return** $(\text{FALSE}, \mathbf{w}_2)$
16: **else If** $\mathbf{w}_2 \supseteq \mathbf{b}$ **then**
17: **Return** $(\text{TRUE}, \mathbf{w}_2)$
18: $\mathbf{w} \leftarrow$ Resolve($\mathbf{w}_1, \mathbf{w}_2$) ▷ Geometric resolution of $\mathbf{w}_1, \mathbf{w}_2$
19: $\mathcal{A} \leftarrow \mathcal{A} \cup \{\mathbf{w}\}$ ▷ Cache the resolution
20: **Return** $(\text{TRUE}, \mathbf{w})$

Next, we explain the resolution step. Resolve is geometric resolution as defined in Section 4.1. Note that by the time Resolve is called in line 18 we know none of $\mathbf{w}_1$ and $\mathbf{w}_2$ covers $\mathbf{b}$. There are a lot of boxes we can infer from $\mathbf{w}_1$ and $\mathbf{w}_2$ if those two boxes are general dyadic boxes that can overlap in peculiar ways. However, TetrisSkeleton forces $\mathbf{w}_1$ and $\mathbf{w}_2$ to be somewhat special, making resolution much more intuitive and clean. In Lemma C.1, we show that all the resolutions in line 18 are ordered geometric resolutions (see Definition 4.2).

We now state the key analytical tool that will be used throughout this paper to bound the runtime of our algorithm in different settings. The tool is a very simple but important combinatorial lemma that says the following: hiding behind the potential poly-log factor in $\tilde{O}$, we can bound the runtime of Tetris by the number of resolutions it performs. The main observation is that in *most* cases when the algorithm backtracks, it does one resolution. The amount of work it does modulo the recursive calls is $\tilde{O}(1)$: inserting a new box, querying for boxes containing a box, and resolving. Finally, line 19 and line 1 make sure that we are not repeating any resolution more than once.

Lemma 4.3 (Runtime is bounded by #resolutions)**.** *Let M denote the total number of resolutions performed by Algorithm 2. Then, the total runtime of Algorithm 2 is $\tilde{O}(M)$.*

4.2.4 Tetris as DPLL with clause learning

We briefly explain how Tetris can be viewed as a form of DPLL with clause learning. A tuple in the output space is an n-dimensional dyadic box each of whose components is a string of length d. When viewed as a bit-string, this tuple is a truth assignment. A dyadic gap box $\mathbf{w}$ under this view can be encoded with a clause, containing all tuples *not* belonging to $\mathbf{w}$. Under this encoding, geometric resolution becomes a particular form of propositional logic resolution. The resolvent of a geometric resolution is a new clause that was inferred and cached in the computation. Hence, Tetris can be cast as a DPLL algorithm for #SAT with a fixed variable ordering and with a particular way of learning new clauses. (It is for #SAT because the algorithm keeps running even after a satisfying assignment is found.

Algorithm 2 Tetris $(\mathcal{B})$

Input: Oracle access to a set of boxes $\mathcal{B}$.
Output: All tuples not covered by any box in $\mathcal{B}$.

1: Initialize$(\mathcal{A})$
2: $(v, \mathbf{w}) \leftarrow$ TetrisSkeleton$(\langle\lambda, \ldots, \lambda\rangle)$
3: **While** $v =$ FALSE **do**
4: $\quad \mathcal{B}' \leftarrow \{\mathbf{b} \in \mathcal{B} \mid \mathbf{b} \supseteq \mathbf{w}\}$ ▷ from the oracle
5: $\quad$ **If** $\mathcal{B}' = \varnothing$ **then**
6: $\quad\quad$ **Report w** as an output tuple
7: $\quad\quad \mathcal{B}' \leftarrow \{\mathbf{w}\}$
8: $\quad \mathcal{A} \leftarrow \mathcal{A} \cup \mathcal{B}'$ ▷ Amend the knowledge base
9: $\quad (v, \mathbf{w}) \leftarrow$ TetrisSkeleton$(\langle\lambda, \ldots, \lambda\rangle)$

See Appendix I.)

4.3 Worst-case Results

The initialization of the knowledge base $\mathcal{A}$ has a crucial implication in terms of the kind of runtime result Tetris is able to attain. In this section, we discuss one extreme where we can load the knowledge base $\mathcal{A}$ with *all* boxes from the input set of boxes $\mathcal{B}$. For notational convenience, we call Tetris with this specific instantiation of Initialize to be Tetris-Preloaded.

It turns out that Tetris-Preloaded achieves the following type of runtime guarantee: given a join query Q, under some assumption about the type of boxes in $\mathcal{B}(Q)$, Tetris-Preloaded runs in time at most the maximum AGM-bound on a bag of any tree decomposition of Q. And we can construct $\mathcal{B}(Q)$ satisfying the assumption in time linear in the input relations' sizes.

Since the AGM-bound result requires some lengthy definitions, we state below a slightly weaker result, in terms of the fractional hypertree width of the query Q. We prove our full (stronger) result in Appendix D.

Theorem 4.4 (Tetris-Preloaded achieves fractional hypertree width bound)**.** *Let Q be a join query, N the total number of input tuples, fhtw the fractional hypertree width of the query, and Z the total number of output tuples. Then, there exists a global attribute order (GAO) σ such that the following holds. Suppose for all $R \in atoms(Q)$, $\mathcal{B}(R)$ is σ-consistent. Then, by setting SAO to be σ,* Tetris-Preloaded *on input $\mathcal{B}(Q)$ runs in time $\tilde{O}(N + N^{fhtw} + Z)$.*

Recall that Tetris uses ordered geometric resolution. It turns out that Tree Ordered Geometric Resolution is enough to recover the AGM bound (see Theorem 5.1). However, Tree Ordered Geometric Resolution is not powerful enough to recover Theorem 4.4 (see Theorem 5.2).

4.4 Beyond Worst-case Results

Our algorithm Tetris not only can recover some existing results as we have seen, but also leads to new results, as presented in this section. In particular, we show that Tetris can extend the bounded treewidth results of [49], which only hold for GAO-consistent input indices, to handle cases of arbitrary input indices, including sophisticated indices such as dyadic trees (and multiple indices per relation).

The crux of *beyond worst-case* guarantee is for the runtime of the algorithm to be measured in the finer notion of (box) certificate size $|\mathcal{C}_\square|$ of the BCP instance, instead of input size. It is easy to construct arbitrarily large input instances for which the certificate size is $\tilde{O}(1)$. (See Appendix B.) Consequently, preloading the knowledge base $\mathcal{A}$ with all boxes from $\mathcal{B}$ as we did with Tetris-Preloaded is no longer an option.

To obtain certificate-based results, we can only load the boxes from $\mathcal{B}$ into $\mathcal{A}$ that are absolutely needed. In particular, we go the other extreme and set $\mathcal{A} \leftarrow \varnothing$ in Initialize$(\mathcal{A})$ (and let lines 4 and 8 in Algorithm 2 load the required boxes from $\mathcal{B}$ into $\mathcal{A}$). For notational convenience, we use
Tetris-Reloaded to refer to Tetris with this specific instantiation of Initialize.

In the following results, we use the well-known fact that if a hypergraph (or a query) has treewidth w, then there is a vertex ordering (or an attribute ordering) that has an *elimination width* w; and, this ordering can be computed in $\tilde{O}(1)$-time in data complexity. We get a near-optimal result for treewidth-1 queries:

Theorem 4.5 ($\tilde{O}(|\mathcal{C}_\square| + Z)$-runtime for treewidth 1)**.** *For any set of boxes $\mathcal{B}$ with $\mathsf{tw}(\mathcal{H}(\mathcal{B})) = 1$, by setting SAO to be the attribute ordering with elimination width 1,* Tetris-Reloaded *solves* BCP *on input $\mathcal{B}$ in time $\tilde{O}(|\mathcal{C}_\square| + Z)$.*

Along with Propositions 3.9 and 3.6, the above result implies the following:

Corollary 4.6. Tetris-Reloaded *evaluates any join query Q with treewidth 1 in time $\tilde{O}(|\mathcal{C}_\square| + Z)$.*

Note that a treewidth of 1 implies that all relations are binary. In Proposition G.12 in Appendix G, we show that as soon as there is a relation of arity $\geqslant 3$, a runtime of $\tilde{O}(|\mathcal{C}_\square| + Z)$ is not possible modulo the hardness of 3SUM.

For general treewidths, we prove a slightly weaker result.

Theorem 4.7 ($\tilde{O}(|\mathcal{C}_\square|^{w+1} + Z)$-runtime for treewidth w)**.** *For any set of boxes $\mathcal{B}$ with $\mathsf{tw}(\mathcal{H}(\mathcal{B})) = w$, by setting SAO to be the attribute ordering with elimination width w,* Tetris-Reloaded *solves* BCP *on input $\mathcal{B}$ in time $\tilde{O}(|\mathcal{C}_\square|^{w+1} + Z)$.*

Along with Propositions 3.9 and 3.6, the above result implies the following:

Corollary 4.8. Tetris-Reloaded *evaluates any join query Q with treewidth w in time $\tilde{O}(|\mathcal{C}_\square|^{w+1} + Z)$.*

4.5 Arbitrary queries

We now show that an enhancement of Tetris gives an improved beyond worst-case result for arbitrary join queries. In particular, we show the following result (see Corollary F.14 in Appendix F.7 for a more general result):

Theorem 4.9. *For any integer $n \geqslant 2$, the problem* BCP *on n dimensions can be solved in time $\tilde{O}(|\mathcal{C}_\square|^{n/2} + Z)$.*

Theorem 5.4 shows that this result cannot be achieved by an algorithm that only performs ordered geometric resolution (like Tetris), no matter which SAO it chooses. The main reason is that it might get stuck in resolving boxes along a fixed dimension due to the fixed SAO, while it could have covered the entire space faster by dynamically switching to resolutions in other dimensions.

We get around this bottleneck by transforming the input boxes into boxes in a higher-dimensional space, then applying Tetris. The idea is to carefully construct this map so that the amount of work per dimension is balanced out. It is worth noting that we are still using the same algorithm Tetris, under a transformed input. Since the analysis of the algorithm is quite involved, we sketch in this section some of the key ideas by making some assumptions about the input. In particular, in this section we assume that the algorithm is given as its input the box certificate of the instance of the BCP, which we will denote by $\mathcal{C}$. We will also refer to this version of the BCP as the *offline* case of the problem. At the end of the section, we outline how we can remove this restriction, leaving the full description to Appendix F.6.

4.5.1 Divide and conquer

To build intuition, we start off with a very special case. Call the input box set $\mathcal{C}$ *balanced* if there exists an attribute X and a partition P_X of the domain $\mathbf{D}(X)$ into $\tilde{O}(\sqrt{|\mathcal{C}|})$ many (disjoint) dyadic intervals such that (i) for each box $\mathbf{b} \in \mathcal{C}$, the interval $\pi_X(\mathbf{b})$ is contained in one of the intervals in P_X and (ii) for each interval $x \in P_X$, the number of boxes $\mathbf{b} \in \mathcal{C}$ such that $\pi_X(\mathbf{b}) \subseteq x$ is also bounded by $\tilde{O}(\sqrt{|\mathcal{C}|})$. The idea is that, when $\mathcal{C}$ is balanced we can solve $\tilde{O}(\sqrt{|\mathcal{C}|})$ independent sub-instances of BCP, one for each *layer* $x \in P_X$; each sub-instance has an input box set of size $\tilde{O}(\sqrt{|\mathcal{C}|})$. This divide and conquer strategy is useful because Tetris can solve BCP on $\mathcal{B}$ in time $\tilde{O}(|\mathcal{B}|^{n-1} + Z)$ (Theorem E.10 in Appendix E). This means if we apply Tetris to each of the $\tilde{O}(\sqrt{|\mathcal{C}|})$ independent subproblems, and then output the union of the results, then we have an overall run time of $\tilde{O}(\sqrt{|\mathcal{C}|} \cdot (\sqrt{|\mathcal{C}|})^{n-1} + Z) = \tilde{O}(|\mathcal{C}|^{n/2} + Z)$, as desired.

In general, the above is too strong a condition. Appendix F.8 gives an explicit certificate $\mathcal{C}$ that is not balanced.

To rectify this situation, we perform a conceptually simple pre-processing step. We design a procedure called Balance that takes as input the certificate $\mathcal{C}$ (recall that we are in the offline case). It outputs a box set $\mathcal{C}'$ of the same size. $\mathcal{C}'$ has a specific SAO such that if one runs Tetris on $\mathcal{C}'$ with this SAO, then one ends up with the desired $\tilde{O}(|\mathcal{C}|^{n/2} + Z)$ runtime. Thus, armed with the balancing procedure, our final algorithm has a very simple structure as illustrated in Algorithm 3.

Algorithm 3 Tetris-Preloaded-LB

Input: A set of boxes C
Output: Output tuples for the BCP on C
1: $\mathcal{B} \leftarrow$ Balance(C)
2: **Return** Tetris-Preloaded($\mathcal{B}$)

4.5.2 Load-balancing with Balance

To sketch out how load-balancing works, we make one further simplifying assumption that we are trying to solve only the Boolean version of BCP: given the set of boxes C, does the union of the boxes in C cover the entire output space? We begin by formalizing certain notions that we used in defining a balanced C above.

Definition 4.10 (Dimension partition)**.** A *partition* P of $\mathbf{D} = \{0,1\}^d$ is a collection of disjoint dyadic intervals whose union is exactly $\mathbf{D}$. Given a dimension X of the BCP, an *X-partition* is a partition of the domain $\mathbf{D}(X)$. We will typically use P_X to denote a partition along dimension X.

Geometrically, a partition along dimension X divides the output space into $|P_X|$ layers, one for each interval x in P_X. An input gap box whose X-component is disjoint from x will not affect whether the x-layer is covered. Hence, to verify whether the x-layer is covered, we can ignore all gap boxes that do not intersect x. If the remaining set of gap boxes is small, then this verification is fast. At the same time, we do not want too many layers because that certainly increases the total amount of verification work. This balancing act leads to our first idea: we find a dimension partition that is somehow balanced.

More concretely, given a set of boxes C and a dyadic interval x on the domain $\mathbf{D}(X)$, define two sets:

$$C_{\subset x}(X) = \{\mathbf{b} \in C \mid \pi_X(\mathbf{b}) \subsetneq x\} \tag{4}$$
$$C_{\supseteq x}(X) = \{\mathbf{b} \in C \mid \pi_X(\mathbf{b}) \supseteq x\}. \tag{5}$$

Note that, for every box $\mathbf{b} \in C \backslash (C_{\subset x}(X) \cup C_{\supseteq x}(X))$, the dyadic interval $\pi_X(\mathbf{b})$ is completely disjoint from the interval x.

Definition 4.11 (Balanced dimension partition)**.** Let C be the set of input gap boxes, and X be any attribute. A *balanced X-partition* is an X-partition P_X such that

$$|P_X| = \tilde{O}(\sqrt{|C|})$$
$$|C_{\subset x}(X)| \leqslant \sqrt{|C|}, \text{ for every } x \in P_X.$$

Given a set C of input gap boxes and an arbitrary attribute X, we can show that a balanced X-partition can be computed in time $\tilde{O}(|C|)$. (Proposition F.4 in Appendix F.) Furthermore, if there exists a balanced X-partition (for some dimension X) such that for every $x \in P_X$, we also have $|C_{\supseteq x}(X)| \leqslant \tilde{O}(\sqrt{|C|})$ and $C_{\supseteq x}(X) \cap C_{\supseteq x'}(X) = \varnothing$ for every $x \neq x' \in P_X$, then C is balanced. So, to have a balanced partition, two fairly strong conditions are required. We next introduce two ideas to relax these requirements.

First, the requirement that $C_{\supseteq x}(X)$ be disjoint for distinct $x \in P_X$ is not strictly required for the divide and conquer strategy to go through. In particular, for each $x \in P_X$ we can create the sub-instance of BCP by copying every box $\mathbf{b} \in C_{\supseteq x}(X)$ and replacing $\pi_X(\mathbf{b})$ by x (for each copy). These new boxes along with $C_{\subset x}(X)$ form a sub-instance. Now we can solve each sub-instance separately as before. Now, if $|C_{\supseteq x}(X)| \leqslant \tilde{O}(\sqrt{|C|})$ for every $x \in P_X$, then we would be done.

Second, unfortunately in general $\left|C_{\supseteq x}(X)\right|$ could be as large as $\Omega(|C|)$. The copying trick to divide the BCP instance into disjoint sub-instances for each $x \in P_X$ is too expensive when $n > 3$. However, when $n = 3$, note that we have $|P_X|$ disjoint BCP instances on dimension 2. In two dimensions, thanks to Theorem E.10, we know Tetris can solve the sub-instance of BCP in time $\tilde{O}(|C|)$, which would lead to the desired $\tilde{O}(|C|^{3/2})$ time (since $|P_X| \leqslant \tilde{O}(\sqrt{|C|})$). Applying bruteforcely, this trick will lead to an overall runtime of $\tilde{O}(|C|^{n-3/2})$, which matches our desired bound of $\tilde{O}(|C|^{n/2})$ only for $n = 3$.

For general n we use the following natural recursive strategy. We divide up the original BCP on n dimensions to $|P_X|$ many disjoint BCP problems (with the copying trick as above) on $n-1$ dimensions (i.e. on all dimensions except X). If a sub-problem has size $\tilde{O}(\sqrt{|C|})$ then we can run Tetris directly. Otherwise we recurse on each of the sub-problems. However, there are two technical issues we need to solve to properly implement this recursive strategy.

First, in the discussion above, we did not talk about boxes in $C_{\subset x}(X)$ for $x \in P_X$. In particular, to define the disjoint $(n-1)$-dimensional sub-problems, we have to perform all possible resolutions of boxes in $C_{\subset x}(X)$ and only retain those boxes $\mathbf{b}$ such that $\pi_X(\mathbf{b}) = x$ for the sub-problem corresponding to $x \in P_X$. Theorem E.10 implies that the number of such boxes could potentially be as large as $\tilde{O}(|C|^{(n-1)/2})$. When $n = 3$, this is still not a problem. However when $n > 3$, we will need to reason about such boxes carefully. (In particular, we cannot consider a sub-problem on $\tilde{O}(|C|^{(n-1)/2})$ many boxes.)

Second, we also need to be careful about the number of times we apply the copying trick above. Some of the boxes that need to be copied might themselves be previous copies, in which case the copying effect would be accumulated. In particular, it is possible for the outcome of a resolution (call it $\mathbf{b}$) among boxes in $C_{\subset x}(X)$ for $x \in P_X$ to become a box in $C_{\supseteq y}(Y)$ for $y \in P_Y$, where Y is some dimension that is encountered later in the recursion. Now in this case, we have to be careful when we copy $\mathbf{b}$ while applying the copying trick for dimension Y.

To tackle the issues above, our final implementation of Balance ends up mapping the BCP on n dimension to a BCP on $2n-2$ dimensions. Note that in this case we cannot use the simple analysis we used for balanced certificates earlier: since even if we can get $\sqrt{|C|}$ disjoint problems with certificate size $\sqrt{|C|}$, each sub-problem will be in $2n-3$ dimensions. Tetris on each sub-problem will take time $\tilde{O}(|C|^{n-2})$, which is generally too costly. Thus, for our Balance function, we have to do a more careful analysis. We would like to stress that Algorithm 3 is still valid: just that the earlier intuitive analysis needs to be tweaked a fair bit.

We present the full analysis of Algorithm 3 along with the complete definition of Balance in Appendix F. The analysis is a bit involved since we have to carefully analyze the number of input boxes each witness depends on. As alluded to earlier, resolvents from earlier levels of recursive calls can interact with resolutions at lower levels of recursions, which foils a straightforward recursive analysis. However, we prove a recursive structural lemma on how resolvents are supported by appropriate number of boxes from previous resolutions, which is enough to appropriately bound the total number of resolutions.

For the *online* version of the problem where the certificate is not given as input, we use the same strategy as Tetris-Reloaded: load boxes from the input $\mathcal{B}$ only when necessary. The number of boxes loaded is $\tilde{O}(|C|)$. Since we are now loading boxes on the go, the notion of a balanced set of boxes changes over time as new boxes are added. For example, a dyadic interval $x \in P_X$ might define a *good* layer whose sub-problem can be solved efficiently for a while, but as new boxes come the layer might eventually become overloaded. Furthermore, the notion of balancedness depends on the total number of boxes. Thus, a *bad* layer might also become good after some time if new boxes do not intersect this layer. We take care of the above key issues by periodically re-adjusting the partitions and show that the total amount of readjustments is not too high.

5 Lower Bounds and Extensions

In this section, we clarify the classes of geometric resolution that are needed to compute various classes of BCP (and hence joins) and prove their limitations. Then, we prove some conditional lower bounds showing that the restrictions in our beyond worst-case results are necessary. Finally, we present some extensions where we prove sharper upper bounds that depend more on the query structure (but only hold for weaker forms of certificates).

5.1 Limitations of Resolution Strategies

So far in this paper, we have seen the class of geometric resolution in Section 4.5. This is the most general class of resolution we will use in this paper. Recall from Section 1 that we denote this class of resolution by Geometric Resolution. In Sections 4.3 and 4.4, we saw a subclass of Geometric Resolution: ordered geometric resolutions. Recall from Section 1 that we denote this class of resolutions by Ordered Geometric Resolution. It turns out that another subclass of Ordered Geometric Resolution, which we call Tree Ordered Geometric Resolution is also an interesting class. Tree Ordered Geometric Resolution (as mentioned in Section 1) is the subclass of Ordered

Geometric Resolution that only re-uses the input gap boxes: in other words, if an intermediate box has to be used more than once, then all the set of resolutions leading up to the intermediate box has to be repeated.[9]

Figure 2 summarizes the power and limitations of the three classes of resolution above. All of our lower bounds (which we present next) follow by constructing explicit hard examples with an empty output for various classes of resolution that we consider in this paper. (Note that for these hard examples any resolution strategy will have to generate the box $\langle \lambda, \ldots, \lambda \rangle$.)

We begin with the power of Tree Ordered Geometric Resolution. We show that one can modify Tetris so that no resolution results are ever cached (this essentially corresponds to running Algorithm 2 but without line 19 in Algorithm 1) so that one can achieve the AGM bound. Note that this change implies that the modified algorithm falls under Tree Ordered Geometric Resolution. This implies that (see Corollary D.3):

Theorem 5.1 (Tree Ordered Geometric Resolution achieves AGM bound)**.** *Let Q be a join query, N the total number of input tuples, and $\mathsf{AGM}(Q)$ the best AGM bound for this instance. Then there exists a scheme in* Tree Ordered Geometric Resolution *that computes Q with $\tilde{O}(\mathsf{AGM}(Q))$ many resolutions.*

Now recall that Tetris uses Ordered Geometric Resolution and in particular, by Theorem 4.4 Tetris is powerful enough to recover the fractional hypertreewidth bound. In turn, this implies that Ordered Geometric Resolution is enough to compute the BCP on boxes with treewidth 1. Next, we argue that Tree Ordered Geometric Resolution is not powerful enough to recover such a result. (See Theorem G.1 and its proof in the appendix.)

Theorem 5.2. *There exists a query Q with treewidth 1 such that every* Tree Ordered Geometric Resolution *algorithm on input $\mathcal{B}(Q)$ needs to make $\Omega(N^{n/2})$ many resolutions, where N is the number of input tuples.*

We now move to Ordered Geometric Resolution. Since Tetris only uses Ordered Geometric Resolution, Theorem 4.7 immediately implies that there exists an Ordered Geometric Resolution algorithm that can solve the BCP on boxes with treewidth w with $\tilde{O}(|\mathcal{C}_\Box|^{w+1} + Z)$ many resolutions. Next, we show that this in general is the best possible (see Theorem G.6 and its proof).

Theorem 5.3. *There exists a set $\mathcal{B}$ of boxes with $1 < \mathsf{tw}(\mathcal{H}(\mathcal{B})) < n-1$ such that any* Ordered Geometric Resolution *algorithm that solves the* BCP *on $\mathcal{B}$ needs to make $\Omega(|\mathcal{C}_\Box(\mathcal{B})|^{w+1})$ many resolutions.*

We have already seen that we can prove bounds of the form $\tilde{O}(|\mathcal{C}_\Box|^w + Z)$ for the special case of $w = 1$ (Theorem 4.5) and $w = n - 1$ (Theorem E.10). Next, we show the upper bound for $w = n - 1$ is the best possible. (See Theorem G.4 and its proof.)

Theorem 5.4. *There exists a set $\mathcal{B}$ of boxes on n dimensions such that any* Ordered Geometric Resolution *algorithm that solves the* BCP *on $\mathcal{B}$ needs to make $\Omega(|\mathcal{C}_\Box(\mathcal{B})|^{n-1})$ many resolutions.*

Note that the above implies that our move to Geometric Resolution to obtain the bound of $\tilde{O}(|\mathcal{C}_\Box|^{n/2} + Z)$ for the BCP problem on dimension n was necessary. It is natural to wonder if this bound can be further improved. We show that this is not possible with Geometric Resolution. (See Theorem G.10 and its proof.)

Theorem 5.5. *For every $n \geqslant 3$, there exists an instance for the* BCP *on n dimensions on which every* Geometric Resolution *algorithm needs to make $\Omega(|\mathcal{C}_\Box|^{n/2})$ many resolutions.*

The proof of Theorem 5.5 follows by a volume argument. We construct the boxes so that resolving any two of them results in a box with a small volume in the output tuples space. Thus, to cover the box $\langle \lambda, \ldots, \lambda \rangle$ one needs to perform a lot of geometric resolutions.

Another natural question is whether Geometric Resolution is strictly less powerful than general resolution. In Appendix J, we show that this indeed is the case by showing that general resolution can solve the hard instance for the proof of Theorem 5.5 with $\tilde{O}(|\mathcal{C}_\Box|)$ many general resolutions. However, we do not know if general resolution can solve all BCP instances on n dimensions with $o(|\mathcal{C}_\Box|^{n/2})$ resolutions.

[9]The qualifier Tree comes from the following fact. We can consider any set of resolutions in Geometric Resolution (and hence Ordered Geometric Resolution) as a DAG– each box is a node and the inputs to a resolution point towards the output of the resolution. Tree Ordered Geometric Resolution is the subset of Ordered Geometric Resolution, where the resolution DAG is a tree.

5.2 Other Results

A natural question we have not addressed so far is whether we can extend Theorem 4.5 to all β-acyclic queries. These queries do admit linear time algorithms but for the weaker comparison based certificate [49]. We show that under the 3SUM conjecture, one cannot hope for such a result for box certificates for β-acyclic queries if relations are allowed arities of 3: see Proposition G.12 in the appendix.

Finally, we are able to prove upper bounds with better dependence on the query than the result in Section 4.5 *if* we work with weaker notions of certificates. See Appendix H.

6 Conclusion

We presented a simple geometric resolution system that allowed us to derive algorithms that match the efficiency of several of the best known algorithms for worst-case analysis and to derive new results for beyond worst-case analysis. Of purely conceptual interest, these rederivations in our simple framework unify and -we argue- simplify their presentation. More technically, our notion of certificate supports a wide range of indexing schemes, compared to previous work that essentially focused on Btrees with a total attribute order. We are excited about further opportunities to more carefully study the impact of indexing on query performance. In addition, we made a connection to proof complexity via geometric resolution that we believe may further strengthen the connection between constraint satisfaction and database join processing. We conclude with two technical questions. First, as observed by [49] it is not possible to obtain a certificate based result with the fhtw in the exponent. It is a very interesting question to figure out the 'correct' notion of fractional cover for certificate-based results. Second, it would be interesting to extend the results of [53] to the certificate setting.

Acknowledgments

We thank Paul Beame for clarifying the relation of our notion of geometric resolution with general resolution and we thank Javiel Rojas-Ledesma for bringing Klee's measure problem to our attention.

MAK's research is supported in part by NSF grant CCF-1161196. HQN's research is supported in part by NSF grants CNF-1409551 and CCF-1319402. CR gratefully acknowledges the support of DARPA's XDATA Program under No. FA8750-12-2-0335, DEFT Program under No. FA8750-13-2-0039, DARPA's MEMEX program under No. FA8750-14-2-0240, NSF CAREER Award No. IIS-1353606, CCF-1356918 and EarthCube Award under No. ACI-1343760, the ONR N000141210041 and N000141310129, the Sloan Research Fellowship, the Moore Foundation Data Driven Investigator award, and gifts from American Family Insurance, Google, Lightspeed Ventures, and Toshiba. AR's research is supported in by part by NSF grants CCF-0844796 and CCF-1319402.

References

[1] Abiteboul, S., Hull, R., and Vianu, V. *Foundations of Databases*. Addison-Wesley, 1995.

[2] Afshani, P., Barbay, J., and Chan, T. M. Instance-optimal geometric algorithms. In *FOCS* (2009), pp. 129–138.

[3] Allender, E., Hellerstein, L., McCabe, P., Pitassi, T., and Saks, M. E. Minimizing DNF formulas and AC^0_d circuits given a truth table. In *IEEE Conference on Computational Complexity* (2006), IEEE Computer Society, pp. 237–251.

[4] Alon, N. On the number of subgraphs of prescribed type of graphs with a given number of edges. *Israel J. Math. 38*, 1-2 (1981), 116–130.

[5] Arnborg, S., and Proskurowski, A. Linear time algorithms for NP-hard problems restricted to partial k-trees. *Discrete Appl. Math. 23*, 1 (1989), 11–24.

[6] Atserias, A., Grohe, M., and Marx, D. Size bounds and query plans for relational joins. In *FOCS* (2008), IEEE Computer Society, pp. 739–748.

[7] Barbay, J., and Kenyon, C. Adaptive intersection and t-threshold problems. In *SODA* (2002), pp. 390–399.

[8] Barbay, J., and Kenyon, C. Alternation and redundancy analysis of the intersection problem. *ACM Transactions on Algorithms 4*, 1 (2008).

[9] Bertino, E., Ooi, B. C., Sacks-Davis, R., Tan, K.-L., Zobel, J., Shidlovsky, B., and Andronico, D. *Indexing Techniques for Advanced Database Systems*. Springer Publishing Company, Incorporated, 2012.

[10] Beyersdorff, O., Galesi, N., and Lauria, M. Parameterized complexity of DPLL search procedures. *ACM Trans. Comput. Log. 14*, 3 (2013), 20.

[11] Bidyuk, B., and Dechter, R. On finding minimal w-cutset. In *UAI* (2004), D. M. Chickering and J. Y. Halpern, Eds., AUAI Press, pp. 43–50.

[12] Blanas, S., Li, Y., and Patel, J. M. Design and evaluation of main memory hash join algorithms for multi-core CPUs. In *SIGMOD* (2011), ACM, pp. 37–48.

[13] Chan, T. Klee's measure problem made easy. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on* (Oct 2013), pp. 410–419.

[14] Chaudhuri, S. An overview of query optimization in relational systems. In *PODS* (1998), ACM, pp. 34–43.

[15] Chekuri, C., and Rajaraman, A. Conjunctive query containment revisited. *Theor. Comput. Sci. 239*, 2 (2000), 211–229.

[16] Chen, J., Lu, S., Sze, S.-H., and Zhang, F. Improved algorithms for path, matching, and packing problems. In *SODA* (2007), N. Bansal, K. Pruhs, and C. Stein, Eds., SIAM, pp. 298–307.

[17] Davis, M., Logemann, G., and Loveland, D. A machine program for theorem-proving. *Comm. ACM 5* (1962), 394–397.

[18] Davis, M., and Putnam, H. A computing procedure for quantification theory. *J. Assoc. Comput. Mach. 7* (1960), 201–215.

[19] de Bruijn, N. G., van Ebbenhorst Tengbergen, C., and Kruyswijk, D. On the set of divisors of a number. *Nieuw Arch. Wiskunde (2) 23* (1951), 191–193.

[20] Dechter, R. Enhancement schemes for constraint processing: Backjumping, learning, and cutset decomposition. *Artif. Intell. 41*, 3 (1990), 273–312.

[21] Dechter, R. *Constraint Processing*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.

[22] Dechter, R., and Pearl, J. Tree-clustering schemes for constraint-processing. In *AAAI* (1988), H. E. Shrobe, T. M. Mitchell, and R. G. Smith, Eds., AAAI Press / The MIT Press, pp. 150–154.

[23] Dechter, R., and Pearl, J. Tree clustering for constraint networks. *Artificial Intelligence 38*, 3 (1989), 353–366.

[24] Dechter, R., and Rish, I. Directional resolution: The Davis-Putnam procedure, revisited. In *KR* (1994), J. Doyle, E. Sandewall, and P. Torasso, Eds., Morgan Kaufmann, pp. 134–145.

[25] Demaine, E. D., López-Ortiz, A., and Munro, J. I. Adaptive set intersections, unions, and differences. In *SODA* (2000), pp. 743–752.

[26] Eén, N., and Sörensson, N. An extensible sat-solver. In *SAT* (2003), E. Giunchiglia and A. Tacchella, Eds., vol. 2919 of *Lecture Notes in Computer Science*, Springer, pp. 502–518.

[27] Fagin, R. Degrees of acyclicity for hypergraphs and relational database schemes. *J. ACM 30*, 3 (1983), 514–550.

[28] Fagin, R., Lotem, A., and Naor, M. Optimal aggregation algorithms for middleware. In *Proceedings of the Twentieth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems* (New York, NY, USA, 2001), PODS '01, ACM, pp. 102–113.

[29] Friedgut, E., and Kahn, J. On the number of copies of one hypergraph in another. *Israel J. Math. 105* (1998), 251–256.

[30] Goerdt, A. Davis-Putnam resolution versus unrestricted resolution. *Ann. Math. Artificial Intelligence 6*, 1-3 (1992), 169–184.

[31] Gottlob, G., Leone, N., and Scarcello, F. Robbers, marshals, and guards: game theoretic and logical characterizations of hypertree width. *J. Comput. Syst. Sci. 66*, 4 (2003), 775–808.

[32] Graefe, G. Query evaluation techniques for large databases. *ACM Computing Surveys 25*, 2 (June 1993), 73–170.

[33] Griggs, J. R. Maximum antichains in the product of chains. *Order 1*, 1 (1984), 21–28.

[34] Grohe, M. Bounds and algorithms for joins via fractional edge covers. In *In Search of Elegance in the Theory and Practice of Computation* (2013), V. Tannen, L. Wong, L. Libkin, W. Fan, W.-C. Tan, and M. P. Fourman, Eds., vol. 8000 of *Lecture Notes in Computer Science*, Springer, pp. 321–338.

[35] Grohe, M., and Marx, D. Constraint solving via fractional edge covers. In *SODA* (2006), ACM Press, pp. 289–298.

[36] Gyssens, M., Jeavons, P., and Cohen, D. A. Decomposing constraint satisfaction problems using database techniques. *Artif. Intell. 66*, 1 (1994), 57–89.

[37] Gyssens, M., and Paredaens, J. A decomposition methodology for cyclic databases. In *Advances in Data Base Theory* (1982), pp. 85–122.

[38] Kim, C., Kaldewey, T., Lee, V. W., Sedlar, E., Nguyen, A. D., Satish, N., Chhugani, J., Di Blas, A., and Dubey, P. Sort vs. hash revisited: fast join implementation on modern multi-core CPUs. *Proc. VLDB Endow. 2*, 2 (Aug. 2009), 1378–1389.

[39] Kolaitis, P. G., and Vardi, M. Y. Conjunctive-query containment and constraint satisfaction. *J. Comput. Syst. Sci. 61*, 2 (2000), 302–332.

[40] Maier, D. *The Theory of Relational Databases*. Computer Science Press, 1983.

[41] Marques-Silva, J., Lynce, I., and Malik, S. Conflict-driven clause learning SAT solvers. *Handbook of satisfiability 185* (2009), 131–153.

[42] Marx, D. Approximating fractional hypertree width. *ACM Transactions on Algorithms 6*, 2 (2010).

[43] Marx, D. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. In *STOC* (2010), pp. 735–744.

[44] Marx, D. Tractable structures for constraint satisfaction with truth tables. *Theory Comput. Syst. 48*, 3 (2011), 444–464.

[45] Marx, D. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM 60*, 6 (2013), 42.

[46] Masek, W. J. Some NP-complete set covering problems. unpublished manuscript.

[47] Milo, R., Shen-Orr, S., Itzkovitz, S., Kashtan, N., Chklovskii, D., and Alon, U. Network motifs: simple building blocks of complex networks. *Science 298*, 5594 (October 2002), 824–827.

[48] Moskewicz, M. W., Madigan, C. F., Zhao, Y., Zhang, L., and Malik, S. Chaff: Engineering an efficient sat solver. In *DAC* (2001), ACM, pp. 530–535.

[49] Ngo, H. Q., Nguyen, D. T., Ré, C., and Rudra, A. Beyond worst-case analysis for joins with Minesweeper. In *PODS* (2014), pp. 234–245.

[50] Ngo, H. Q., Porat, E., Ré, C., and Rudra, A. Worst-case optimal join algorithms: [extended abstract]. In *PODS* (2012), pp. 37–48.

[51] Ngo, H. Q., Ré, C., and Rudra, A. Skew strikes back: New developments in the theory of join algorithms. In *SIGMOD RECORD* (2013), pp. 5–16.

[52] Nguyen, D., Aref, M., Bravenboer, M., Kollias, G., Ngo, H. Q., Ré, C., and Rudra, A. Join processing for graph patterns: An old dog with new tricks. In *Proceedings of the GRADES'15* (New York, NY, USA, 2015), GRADES'15, ACM, pp. 2:1–2:8.

[53] Olteanu, D., and Zavodny, J. Size bounds for factorised representations of query results. *ACM Transactions on Database Systems* (2014). To appear.

[54] O'Neil, P. E., and Graefe, G. Multi-table joins through bitmapped join indices. *SIGMOD Record 24*, 3 (1995), 8–11.

[55] O'Neil, P. E., and Quass, D. Improved query performance with variant indexes. In *SIGMOD Conference* (1997), J. Peckham, Ed., ACM Press, pp. 38–49.

[56] Overmars, M. H., and Yap, C.-K. New upper bounds in klees measure problem. *SIAM Journal on Computing 20*, 6 (1991), 1034–1045.

[57] Papadimitriou, C. H., and Yannakakis, M. On the complexity of database queries. In *PODS* (1997), pp. 12–19.

[58] Pătraşcu, M. Towards polynomial lower bounds for dynamic problems. In *Proc. 42nd ACM Symposium on Theory of Computing (STOC)* (2010), pp. 603–610.

[59] Pearl, J. *Probabilistic reasoning in intelligent systems - networks of plausible inference*. Morgan Kaufmann series in representation and reasoning. Morgan Kaufmann, 1989.

[60] Przulj, N., Corneil, D. G., and Jurisica, I. Modeling interactome: scale-free or geometric? *Bioinformatics 20*, 18 (2004), 3508–3515.

[61] Ramakrishnan, R., and Gehrke, J. *Database Management Systems*, 3 ed. McGraw-Hill, Inc., New York, NY, USA, 2003.

[62] Robertson, N., and Seymour, P. D. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms 7*, 3 (1986), 309–322.

[63] Roughgarden, T. Lecture notes on beyond worst-case analysis (cs264).

[64] Scarcello, F. Query answering exploiting structural properties. *SIGMOD Record 34*, 3 (2005), 91–99.

[65] Silva, J. P. M., and Sakallah, K. A. Grasp - a new search algorithm for satisfiability. In *ICCAD* (1996), pp. 220–227.

[66] Silva, J. P. M., and Sakallah, K. A. Grasp: A search algorithm for propositional satisfiability. *IEEE Trans. Computers 48*, 5 (1999), 506–521.

[67] Spiegler, I., and Maayan, R. Storage and retrieval considerations of binary data bases. *Inf. Process. Manage. 21*, 3 (1985), 233–254.

[68] Suri, S., and Vassilvitskii, S. Counting triangles and the curse of the last reducer. In *WWW* (2011), pp. 607–614.

[69] Tsourakakis, C. E. Fast counting of triangles in large real networks without counting: Algorithms and laws. In *ICDM* (2008), IEEE Computer Society, pp. 608–617.

[70] Ullman, J. D. *Principles of Database and Knowledge-Base Systems, Volume II*. Computer Science Press, 1989.

[71] Veldhuizen, T. L. Triejoin: A simple, worst-case optimal join algorithm. In *ICDT* (2014), pp. 96–106.

[72] Yannakakis, M. Algorithms for acyclic database schemes. In *VLDB* (1981), pp. 82–94.

[73] Zhang, H. Sato: An efficient propositional prover. In *CADE* (1997), W. McCune, Ed., vol. 1249 of *Lecture Notes in Computer Science*, Springer, pp. 272–275.

A Background

A.1 Join query

Let $\mathcal{A}$ be a set of attribute names, where an attribute $A \in \mathcal{A}$ is a variable over a finite and discrete domain $\mathbf{D}(A)$. Let $\mathcal{R}$ be a set of *relation symbols*. A *relational schema* for the symbol $R \in \mathcal{R}$ of arity k is a tuple $\text{vars}(R) = (A_{i_1}, \ldots, A_{i_k})$ of distinct attributes that defines the attributes of the relation. A relational database schema is a set of relational symbols and associated schemas denoted by $R(\text{vars}(R)), R \in \mathcal{R}$. A relational instance for $R(A_{i_1}, \ldots, A_{i_k})$ is a subset of $\mathbf{D}(A_{i_1}) \times \cdots \times \mathbf{D}(A_{i_k})$. A relational database $\mathcal{D}$ is a collection of instances, one for each relational symbol in the schema, denoted by $R^{\mathcal{D}}$. Often the database is clear from context and we drop the superscript $\mathcal{D}$ from the relation symbols.

A *natural join* query (or simply join query) Q is specified by a finite subset of relational symbols $\text{atoms}(Q) \subseteq \mathcal{R}$, denoted by $\bowtie_{R \in \text{atoms}(Q)} R$. Let $\text{vars}(Q)$ denote the set of all attributes that appear in some relation in Q, that is

$$\text{vars}(Q) = \{A \mid A \in \text{vars}(R) \text{ for some } R \in \text{atoms}(Q)\}.$$

Given a database instance $\mathcal{D}$, the output of the query Q on the database instance $\mathcal{D}$ is denoted $Q(\mathcal{D})$ and is defined as

$$Q(\mathcal{D}) \stackrel{\text{def}}{=} \left\{\mathbf{t} \in \mathbf{D}^{\text{vars}(Q)} \mid \pi_{\text{vars}(R)}(\mathbf{t}) \in R^{\mathcal{D}} \text{ for each } R \in \text{atoms}(Q)\right\}$$

where $\mathbf{D}^{\text{vars}(Q)}$ is a shorthand for $\times_{A \in \text{vars}(Q)} \mathbf{D}(A)$, and π is the projection operator. When the instance is clear from the context we will refer to $Q(\mathcal{D})$ by just Q.

For example, in the following so-called *triangle query*

$$Q_\triangle = R(A, B) \bowtie S(B, C) \bowtie T(A, C).$$

we have $\text{vars}(Q_\triangle) = \{A, B, C\}$, $\text{vars}(R) = \{A, B\}$, $\text{vars}(S) = \{B, C\}$, and $\text{vars}(T) = \{A, C\}$.

A.2 AGM bound

The structure of a join query Q can be represented by a hypergraph $\mathcal{H}(Q)$, or simply $\mathcal{H}$. The hypergraph has vertex set $\mathcal{V} = \text{vars}(Q)$, and edge set

$$\mathcal{E} = \{\text{vars}(R) \mid R \in \text{atoms}(Q)\}.$$

We often index the relations using edges from this hypergraph. Hence, instead of writing $R(\text{vars}(R))$, we can write R_F, for $F \in \mathcal{E}$.

A *fractional edge cover* of a hypergraph $\mathcal{H}$ is a point $\mathbf{x} = (x_F)_{F \in \mathcal{E}}$ in the following polyhedron:

$$\left\{\mathbf{x} \mid \sum_{F : v \in F} x_F \geqslant 1, \forall v \in \mathcal{V}, \mathbf{x} \geqslant \mathbf{0}\right\}.$$

Atserias-Grohe-Marx [6] and Grohe-Marx [35] proved the following remarkable inequality, which shall be referred to as the *AGM's inequality* henceforth. For any fractional edge cover $\mathbf{x}$ of the query's hypergraph,

$$|Q| = |\bowtie_{F \in \mathcal{E}} R_F| \leqslant \prod_{F \in \mathcal{E}} |R_F|^{x_F}. \tag{6}$$

Here, $|Q|$ is the number of tuples in the (output) relation Q.

The optimal edge cover for the AGM bound depends on the relation sizes. To minimize the right hand side of (6), we can solve the following linear program:

$$\begin{aligned} \min \quad & \sum_{F \in \mathcal{E}} (\log_2 |R_F|) \cdot x_F \\ \text{s.t.} \quad & \sum_{F : v \in F} x_F \geqslant 1, v \in \mathcal{V} \\ & \mathbf{x} \geqslant \mathbf{0} \end{aligned}$$

Implicitly, the objective function above depends on the database instance $\mathcal{D}$ on which the query is applied. Let $\rho^*(Q,\mathcal{D})$ denote the optimal objective value to the above linear program. We refer to $\rho^*(Q,\mathcal{D})$ as the *fractional edge cover number* of the query Q with respect to the database instance $\mathcal{D}$, following Grohe [34]. The AGM's inequality can be summarized simply by $|Q| \leqslant 2^{\rho^*(Q,\mathcal{D})}$.

Definition A.1 (AGM(Q))**.** We also define $\mathsf{AGM}(Q,\mathcal{D}) = 2^{\rho^*(Q,\mathcal{D})}$; in particular $\mathsf{AGM}(Q,\mathcal{D})$ is the best AGM-bound for query Q on database instance $\mathcal{D}$.

When $\mathcal{D}$ is clear from context, we will simply use $\rho^*(Q)$ to denote $\rho^*(Q,\mathcal{D})$, and $\mathsf{AGM}(Q)$ to denote $\mathsf{AGM}(Q,\mathcal{D})$.

Definition A.2 (Fractional edge cover number)**.** Sometimes it is convenient to replace individual relation sizes by a single number N representing the entire input size. In that case, in the linear program above we can drop the $(\log_2 |R_F|)$ factors. The optimal objective value for this simplified program is called the *fractional edge cover number* of the hypergraph $\mathcal{H}$, denoted by $\rho^*(\mathcal{H})$.

A.3 Tree decomposition, acyclicity, and various notions of widths

Definition A.3 (Acyclicity and GYO elimination)**.** There are many definitions of acyclic hypergraphs. A hypergraph $(\mathcal{V},\mathcal{E})$ is α-acyclic if the GYO procedure returns an empty hypergraph [1, p. 128]. Essentially, in GYO one iterates two steps: (1) remove any edge that is empty or contained in another hyperedge, or (2) remove vertices that appear in at most one hyperedge. If the result is empty, then the hypergraph is α-acyclic. A hypergraph $\mathcal{H}$ is β-acyclic if the hypergraph formed by any subset of hyperedges of $\mathcal{H}$ is α-acyclic.

Definition A.4 (Tree decomposition)**.** Let $\mathcal{H} = (\mathcal{V},\mathcal{E})$ be a hypergraph. A *tree decomposition* of $\mathcal{H}$ is a pair (T,χ) where $T = (V(T),E(T))$ is a tree and $\chi : V(T) \to 2^{\mathcal{V}}$ assigns to each node of the tree T a set of vertices of $\mathcal{H}$. The sets $\chi(t)$, $t \in V(T)$, are called the *bags* of the tree decomposition. There are two properties the bags must satisfy:

(a) For every hyperedge $F \in \mathcal{E}$, there is a bag $\chi(T)$ such that $F \subseteq \chi(t)$.

(b) For every vertex $v \in \mathcal{V}$, the set $\{t \mid t \in T, v \in \chi(t)\}$ is not empty and forms a connected subtree of T.

The rest of this section roughly follows the definitions given in [42]. The *width* of a tree decomposition is the quantity

$$\max_{t \in V(T)} |\chi(t)| - 1.$$

The *treewidth* of a hypergraph $\mathcal{H}$, denoted by $\mathsf{tw}(\mathcal{H})$, is the minimum width over all tree decompositions of the hypergraph. Let $\rho^*(t)$ denote the fractional edge cover number of the hypergraph $\mathcal{H}[\chi(t)]$: the hypergraph $\mathcal{H}$ restricted to the bag $\chi(t)$. Then the quantify

$$\max_{t \in V(T)} |\rho^*(t)|$$

is called the *fractional hypertree width* (fhtw) of this tree decomposition of the hypergraph $\mathcal{H}$. The fractional hypertree width of $\mathcal{H}$, denoted by $\mathrm{fhtw}(\mathcal{H})$, is the minimum one over all tree decompositions of $\mathcal{H}$. It should be clear that $\mathrm{fhtw}(\mathcal{H}) \leqslant \mathrm{tw}(\mathcal{H})$.

Definition A.5 (Treewidth of a query)**.** The treewidth of a query Q, denoted by $\mathsf{tw}(Q)$, is the treewidth of its hypergraph $\mathcal{H}(Q)$.

B Indices as Gap Collections, Geometric Certificates

In this section, we attempt to make the gap box abstraction more concrete. We will explain what the "boxes" are and why the idea that every database index is a collection of boxes holds for ordered input index data structures such as a sorted list, B-tree, B+-tree, or a trie. Note again that a hash table can be simulated by a search tree within a log-factor; hence this observation holds for hash-based indices as far as the theoretical bounds in this paper are concerned.

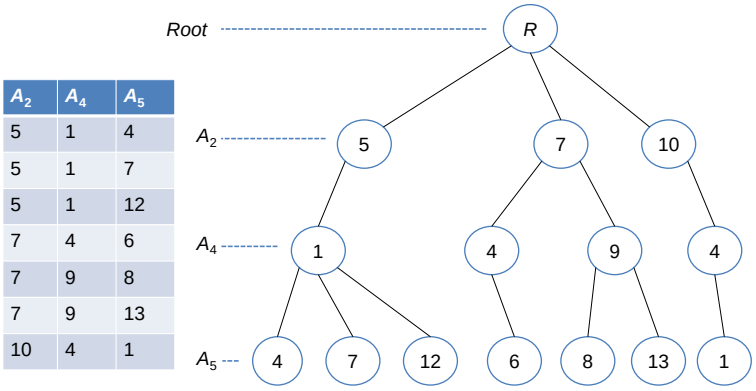

A_2	A_4	A_5
5	1	4
5	1	7
5	1	12
7	4	6
7	9	8
7	9	13
10	4	1

Figure 8: The (unbounded fanout) B-tree data structure used by Minesweeper

B.1 Minesweeper and GAO-consistent certificate

We recall key concepts and ideas from [49] from which a new notion of *geometric certificate* called *GAO-consistent certificate* arises naturally. The Minesweeper algorithm assumes that the input relations are already indexed using a search-tree data structure such as a traditional B-tree which is widely used in commercial relational database systems [61, Ch.10] or a Trie [71]. For example, Figure 8 shows the index for a relation R on attribute set $\text{vars}(R) = \{A_2, A_4, A_5\}$. This index for R is in the order A_2, A_4, A_5. Furthermore, there is an ordering of all the attributes in $\text{vars}(Q)$ – called the *global attribute order* (GAO) – such that all input relations are indexed consistent with this GAO. This assumption shall be referred to as the *GAO-consistency assumption*. For example, for the triangle query $Q = R(A,B) \bowtie S(B,C) \bowtie T(A,C)$, if the GAO is B, A, C, then R is indexed in the (B,A) order, S in the (B,C), and T in the (A,C).

Minesweeper views the set of potential output tuples of a join query as "grid points" in a high-dimensional space called the *output space*. The output space is the cross-product of all domains of input attributes. Minesweeper maintains a data structure called the *constraint data structure* (CDS) which holds a collection of "gaps" in the output space whose union contains all points which are either non-output tuples or output tuples already reported. When the union of gaps in the CDS covers the entire output space, the algorithm terminates.

More concretely, Minesweeper starts off by obtaining an arbitrary "*probe point*" $\mathbf{t}$ from the output space. Given the current probe point, Minesweeper performs the following steps:

1. Minesweeper queries into the indices of the input relations to verify whether the current probe point $\mathbf{t}$ is indeed an output tuple, in which case $\mathbf{t}$ is output and an appropriate gap is inserted into the CDS signifying that the output $\mathbf{t}$ has already been reported. If $\mathbf{t}$ is not an output tuple, the input index structures return some gaps "around" $\mathbf{t}$. These gaps are inserted into the CDS.

2. Then, Minesweeper queries into the CDS to obtain the next probe point, which is a point not covered by the union of the gaps inserted thus far. If no probe point exists, the algorithm terminates, because all inserted gaps cover the output space.

To be more concrete, suppose $\text{vars}(Q) = \{A_1, \ldots, A_6\}$ with $(A_1, A_2, \ldots, A_6)$ being the GAO. Suppose the relation R shown in Figure 8 is an input relation. Consider the following probe point

$$\mathbf{t} = (t_1, t_2, t_3, t_4, t_5, t_6) = (6, 6, 1, 3, 7, 9).$$

We first project this probe point down to the coordinate subspace spanned by the attributes of R: $(t_2, t_4, t_5) = (6, 3, 7)$. From the index structure for R, we see that $t_2 = 6$ falls between the two A_2-values 5 and 7 in the relation. Thus, this index returns a gap consisting of all points lying between the two hyperplanes $A_2 = 5$ and $A_2 = 7$. This gap is encoded with the constraint

$$\langle *, (5,7), *, *, *, * \rangle,$$

where $*$ is the wildcard character matching any value in the corresponding domain, and $(5,7)$ is an open interval on the A_2-axis. On the other hand, suppose the probe point is

$$\mathbf{t} = (t_1, t_2, t_3, t_4, t_5, t_6) = (6, 7, 1, 5, 8, 9).$$

Then, a gap returned might be the band in the hyperplane $A_2 = 7,\ 4 < A_4 < 9$. The encoding of this gap is

$$\langle *, 7, *, (4, 9), *, * \rangle.$$

The number 7 indicates that this gap is inside the hyperplane $A_2 = 7$, and the open interval encodes all points inside this hyperplane where $4 < A_4 < 9$.

Due to the GAO-consistency assumption, all the constraints returned by the input indices have the property that for each constraint there is only one interval component, after that there are only wildcard component. Henceforth, these constraints will be called *GAO-consistent constraints*. The key result from [49] is the following: If we consider the class of join algorithms that only perform comparisons between input elements, then the number of comparisons necessary to certify the correctness of the output is a lowerbound on the run time of this class of (non-deterministic) algorithms. This class of comparison-based algorithms models a very wide class of join algorithms, including index-nested-loop join, block-nested-loop join, hash-join (up to a log-factor), grace join.

Note that every constraint inserted into the CDS is a (multi-dimensional) rectangle inside the output space. We will call these rectangles *boxes*. The analysis from [49] shows the following:

- Let $C^{\text{gao}}_<$ denote the minimum set of comparisons sufficient to certify the output, then the number of probe points $\mathbf{t}$ issued by Minesweeper is $\tilde{O}(|C^{\text{gao}}_<| + Z)$, where Z is the output size, and $\tilde{O}$ hides a query-dependent factor.
- Every probe point $\mathbf{t}$ is a point that is not covered by existing boxes stored in the CDS.
- If a new probe point $\mathbf{t}$ is not an output tuple, then Minesweeper will insert at least one new box containing the probe point $\mathbf{t}$. We will refer to these boxes as *gap boxes*. Note the important fact that Minesweeper often also inserts boxes that *do not* contain $\mathbf{t}$; this is because the analysis needs to show that we can "pay" for this iteration by a fresh comparison in $C^{\text{gao}}_<$.
- If a new probe point $\mathbf{t}$ *is* an output tuple, then Minesweeper will also insert a new box containing $\mathbf{t}$ to rule it out; but we will refer to these boxes as *output boxes*.

Now, let's forget about $C^{\text{gao}}_<$ and examine what the CDS sees and processes. The CDS has a set of *output boxes*, and a collection of *gap boxes* which do not contain any output point. When the CDS cannot find a probe point anymore, the union of output boxes and gap boxes is the entire output space O. In other words, every point in the output space is either an output point, or is covered by a gap box. We will call the collection of gap boxes satisfying this property *box certificate* (to be defined more precisely below): It certifies that the reported output points are all the output points of the join query. A box certificate is a purely geometric notion, and on the surface does not seem to have anything to do with comparisons. Yet from the results in [49], we now know that a box certificate of minimum size is a lowerbound on the number of comparisons issued by *any* comparison-based join algorithm (with the GAO-consistency assumption).

With the GAO-consistency assumption, the gap boxes can only come from the gaps issued by the input relations following their GAO-consistent index structures. We make this notion more precise here.

Definition B.1 (GAO-consistent box certificate)**.** Let Q be a natural join query whose input relations have already been indexed consistent with a GAO $(A_1, \ldots, A_n)$. The gap boxes owned by $R \in \text{atoms}(Q)$ are formed by the gaps between adjacent sibling nodes in the search trees for R. These boxes can be of dimension n, $n-1$, down to 1. (A 1-dimensional box is a segment on a line.) Let $\mathcal{B}(R)$ denote the set of gap boxes owned by R.

A *GAO-consistent box certificate* is a subset of gap boxes from $\bigcup_{R \in \text{atoms}(Q)} \mathcal{B}(R)$ that cover every point not in the output. We use $C^{\text{gao}}_\square$ to also denote an optimal GAO-consistent box certificate, which means it has the minimum number of gap boxes.

Example B.2 (GAO-consistent gap boxes)**.** Consider a relation $S(A, B)$ on two attributes A and B, represented by points shown in Figure 9(a). If S was indexed in the A, B order, then the gap boxes from S (i.e. the set $\mathcal{B}(S)$) are shown in Figure 9(b). If S was indexed in the B, A order, then the gap boxes from S are shown in Figure 9(c).

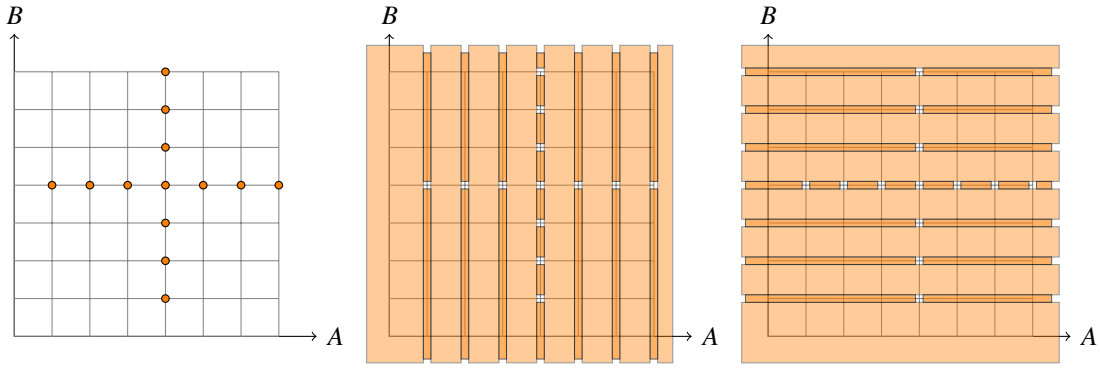

(a) Tuples from $S(A, B)$ (b) (A, B)-consistent gap boxes (c) (B, A)-consistent gap boxes

Figure 9: A relation and its GAO-consistent gap boxes

Example B.3 (GAO-consistent box certificate)**.** Consider the bowtie query

$$Q = R(A) \bowtie S(A, B) \bowtie T(B)$$

where the GAO is (A, B). An input instance to this problem is shown in Figure 10(a). A GAO-consistent box certificate where the GAO is (A, B) is shown in Figure 10(b). A GAO-consistent box certificate where the GAO is (B, A) is shown in Figure 10(c).

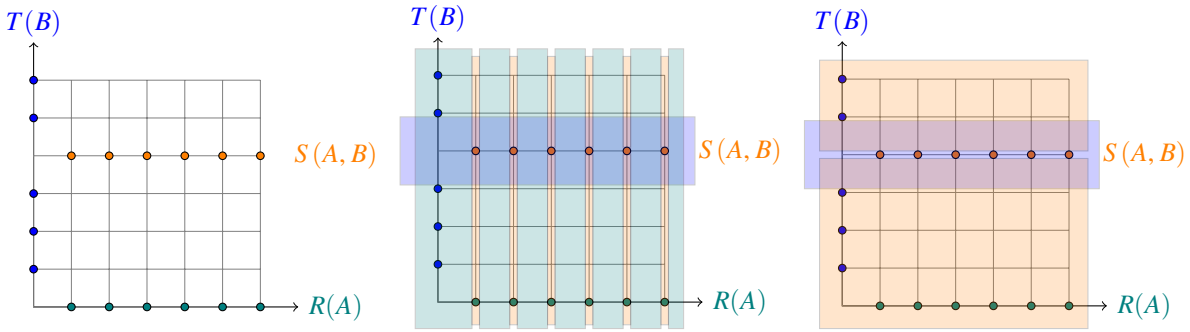

(a) An instance of the bowtie query (b) (A, B)-consistent box certificate (c) (B, A)-consistent box certificate

Figure 10: A GAO-consistent box certificate for the bowtie query

The following proposition follows from the above discussion on the main results of [49], leading the way to a more general geometric notion of certificates.

Proposition B.4 (Implicit from [49])**.** *For any input instance to the join evaluation problem where all input relations are indexed consistent with a fixed GAO, we have* $|C_{\square}^{gao}| = O(|C_{<}^{gao}|)$.

B.2 General box certificates

It can be seen from Example B.3 that a GAO-consistent certificate is highly sensitive to the GAO for the same input data. In the (A, B) GAO, all gap boxes from R are vertical strips, all the gap boxes from S are the same vertical strips as R and the two segments above and below each orange point. The gap boxes from T are horizontal strips between

the blue points. To cover the entire space, we will need $\Omega(N)$ strips. In the (B, A) GAO, S has two big boxes above and below the orange line. So the space can be covered with just three boxes: two orange boxes and one blue box.

A fact often seen in practice [9] is that relations are indexed with multiple search keys. In such case, there is no reason to stick to a particular GAO. For example, if the relation S from Example B.3 was already indexed in *both* the (A, B) order and the (B, A) order, we should make use of the information form both indices to speed up join processing. From this view point, the notion of geometric certificate starts to gain traction: if we stick to a comparison-based algorithm it is not clear how we can make use of the availability of multiple indices per relation. From the geometric view point, we can simply think of the set of indices for every relation as a collection of boxes whose union is exactly the set of tuples *not* in the relation.

Abstracting away, henceforth we will model the input relations as a collection of boxes, denoted by $\mathcal{B}(R)$ for each input relation R. This way, it does not matter anymore how many indices are available per input relation. The notion of *box certificate* comes naturally.

Definition B.5 (Box certificate)**.** Let Q be a natural join query whose input relations have already been indexed. There can be multiple indices per relation. Let R be a k-ary relation on attributes, say, $(A_1, \ldots, A_k)$. The gap boxes owned by $R \in \text{atoms}(Q)$ are a set of rectangles whose union contains precisely the set of tuples on $\mathbf{D}(A_1) \times \cdots \times \mathbf{D}(A_k)$ which *do not* belong to R. Let $\mathcal{B}(R)$ denote the set of gap boxes owned by R.

A *box certificate* is a subset of gap boxes from $\bigcup_{R \in \text{atoms}(Q)} \mathcal{B}(R)$ that cover every point not in the output. We use $C_\square$ to denote a box certificate of minimum size for the instance.

The following proposition is straightforward.

Proposition B.6. *For any input instance to the join problem, and for any fixed GAO, we have* $|C_\square| = O(|C_\square^{gao}|)$ *(and thus* $|C_\square| = O|C_<^{gao}|$*). And, for some input instances* $|C_\square| = O(1)$ *while* $|C_\square^{gao}| = \Omega(N)$.

Proof. The first statement is obvious but needs a bit of clarification. Suppose the input instances are indexed with multiple search key orderings. For example, $S(A, B, C)$ can be indexed with 6 different B-tree key orderings. Then, if we fixed a GAO we are only able to use one of the 6 available indices from S. To see the second statement, we can refer back to Example B.3. If the GAO was (A, B), then obviously $|C_<^{\text{gao}}| = \Omega(N)$, while the general box certificate can take the boxes in Figure 10(c). □

Example B.7 (Gaps from B-tree indices can be bad)**.** A box certificate is a function of the input gap boxes. Perhaps the most natural gap boxes are the gap boxes coming from search tree indices such as B-tree or trie as shown in Figure 9. In this example, we show that even when the input relations are indexed using all possible GAOs, the box certificate can still be much larger than necessary for some input instances.

Consider again the bowtie query $Q = R(A) \bowtie S(A, B) \bowtie T(B)$. The relation $S(A, B)$ is indexed in both directions (A, B) and (B, A). For the input instance shown in Figure 11(a), $\Omega(N)$ rectangles are still required to cover the output space.

To see this, note that all the B-tree-style gap boxes from R and T cannot cover any grid point (i, j) for $i, j \in [7] - \{4\}$. And the gap boxes from S from the $S(A, B)$-index and from the $S(B, A)$-index are shown in Figure 9. The big (light orange) boxes can't cover the grid points. And we will need $\Omega(N)$ thin boxes to cover all of them, in spite of the fact that two gap boxes from R and T already cover all points in S.

Example B.8 (Non-B-tree gap boxes can help)**.** Suppose the indices for relation $S(A, B)$ store four additional gap-boxes as shown in Figure 11(b). Then, we'd have a constant-sized certificate for the bad instance from Figure 11(a). Unfortunately, the commonly used B-tree indices do not allow for returning these types of gap boxes. One of the key contributions of this paper is the observation that the kind of gap boxes shown in Figure 11(b) can be *inferred* from the B-tree gap boxes. In fact, they can be inferred from *any* collection of gap boxes encoding the same relation. The inference framework developed in this paper can be used to analyze the complexity of such inference algorithms, which helps build better input indices.

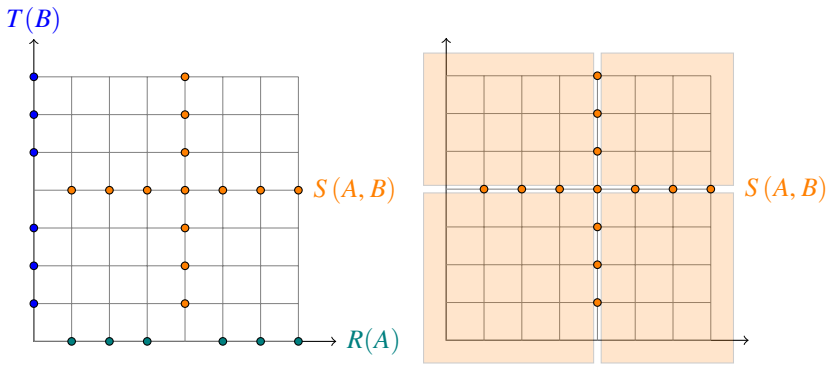

(a) An input relation $S(A, B)$ (b) Additional gap-boxes from S

Figure 11: B-tree gap boxes can still be bad, and more general boxes help

Algorithm 4 Algorithm for discovering a box certificate

```
CDS ← ∅                                                          ▷ No box discovered yet
While CDS can find t not in any stored box do
    If π_{vars(R)}(t) ∈ R for every R ∈ atoms(Q) then
        Report t and insert t as an output box back to the CDS
    else
        Query all R ∈ atoms(Q) for all maximal gap boxes containing π_{vars(R)}(t)
        Insert those gap boxes into CDS
```

B.3 On discovering an optimal box certificate and the case for dyadic box certificate

Algorithm 4 is essentially the skeleton of the Minesweeper algorithm. There are a couple of subtle differences which do not affect our discussion here. If for every input relation R and every probe point $\mathbf{t}$, there are only $\tilde{O}(1)$ maximal gap boxes in $\mathcal{B}(R)$ containing $\pi_{\text{vars}(R)}(\mathbf{t})$, then it is easy to see that Algorithm 4 discovers any box certificate $C_\square$ with $\tilde{O}(m|C_\square| + Z)$ many insertions into CDS, where m is the number of input relations.

For a GAO-consistent index, it certainly is true that there is only *one* maximal GAO-consistent gap box in $\mathcal{B}(R)$ containing $\pi_{\text{vars}(R)}(\mathbf{t})$. However, this is not true in general, when we allow for arbitrary gap boxes. This is because the number of maximal gap-boxes containing a given probe point can be linear in the input size, and it is possible that only one of them is used in the optimal box certificate.

Example B.9 (There can be $\Omega(N)$ maximal gap-boxes)**.** Consider, for example, the situation shown in Figure 12.

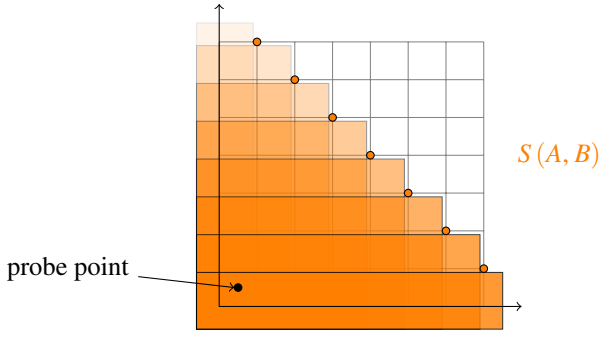

Figure 12: There can be $\Omega(N)$ maximal gap-boxes

In this section, we make the case for a particular kind of gap boxes called *dyadic boxes* to overcome the above problem. We make the following observations (to be proved below):

- for every probe point $\mathbf{t}$, the number of maximal dyadic gap boxes containing $\pi_{\text{vars}(R)}(\mathbf{t})$ is $\tilde{O}(1)$.
- any gap box can be broken down into $\tilde{O}(1)$ dyadic gap boxes.
- for any box certificate $C_\square$, there is a dyadic box certificate of size $\tilde{O}(|C_\square|)$.

From these observations, we will then assume (up to a polylog-factor) in the rest of the paper that the sets $\mathcal{B}(R)$ contain only dyadic gap boxes.

Consider any join query Q with $\text{vars}(Q) = \{A_1, \ldots, A_n\}$. For simplicity, but without any loss of generality, let us assume the domain of each attribute is the set of all binary strings of length d, i.e.

$$\mathbf{D}(A) = \{0,1\}^d, \ \forall A \in \text{vars}(Q).$$

This is equivalent to saying that the domain of each attribute is the set of all integers from 0 to $2^d - 1$. Since d is the number of bits needed to encode a data value of the input, d is logarithmic in the input size.

Definition B.10 (Dyadic interval)**.** A *dyadic interval* is a binary string x of length $\leqslant d$. We use $|x|$ to denote its length. This interval represents all the binary strings y such that $|y| = d$ and x is a prefix of y. Translating to the integral domain, the dyadic interval represents all integers in the range

$$[x2^{d-|x|}, (x+1)2^{d-|x|} - 1].$$

Here, we view x as the integer corresponding to the bit string. The empty string $x = \lambda$ is a dyadic interval consisting of all possible values in the domain. If x has length d, then it is called a *unit dyadic interval* or a *full-length dyadic interval*. In this case, x represents a point in the domain.

Definition B.11 (Dyadic box)**.** Let $\text{vars}(Q) = \{A_1, \ldots, A_n\}$. A *dyadic box* is an n-tuple of dyadic intervals:

$$\mathbf{b} = \langle x_1, \ldots, x_n \rangle,$$

If all components of $\mathbf{b}$ are unit dyadic intervals, then $\mathbf{b}$ represents a *point* in the output space. The dyadic box *is* the set of all tuples

$$\mathbf{t} = (t_1, \ldots, t_n) \in \prod_{i=1}^{n} \mathbf{D}(A_i)$$

such that t_i belongs to the dyadic interval x_i, for all $i \in [n]$.

Note again that some dyadic intervals can be λ, matching arbitrary domain values; also, a dyadic box $\mathbf{b}$ contains a dyadic box $\mathbf{b}'$ if each of $\mathbf{b}$'s components is a prefix of the corresponding component in $\mathbf{b}'$. The set of all dyadic boxes form a partially ordered set (poset) under this containment relation.

Proposition B.12 (Number of (maximal) dyadic boxes containing a point is $\tilde{O}(1)$)**.** *Let R be an arbitrary input relation of arity k. Then, given a probe point $\mathbf{t}$, the number of maximal dyadic gap boxes from R that contain $\pi_{\text{vars}(R)}(\mathbf{t})$ is bounded by*

$$\sum_{i=0}^{k} (-1)^i \binom{k}{i} \binom{\lfloor kd/2 \rfloor + k - 1 - id - i}{k-1} \leqslant \binom{\lfloor kd/2 \rfloor + k - 1}{k-1} \leqslant d^k.$$

The number of dyadic boxes containing $\mathbf{t}$ is at most d^k. (Note that $d^k = \tilde{O}(1)$, because in this paper $\tilde{O}$ hides factors that are query dependent and poly-log dependent on the data.)

Proof. Consider the poset $\mathcal{P}_d = \{0,1\}^{\leqslant d}$ of binary strings of length $\leqslant d$ under the reversed prefix order, i.e. $x \preceq y$ in this poset if y is a prefix of x. (This is just a binary tree in reverse.)

Let $\mathcal{P}_d^k$ denote the kth *Cartesian power* of $\mathcal{P}_d$ (under the product order), i.e. $\mathcal{P}_d^k = \mathcal{P}_d \times \mathcal{P}_d \times \cdots \times \mathcal{P}_d$, k times. Then, $\mathcal{P}_d^k$ is the poset of all dyadic boxes under the containment order. The set of all dyadic boxes *containing* a point $\mathbf{p}$ is a principal filter of $\mathcal{P}_d^k$ at $\mathbf{p}$, denoted by $\mathcal{P}_d^k[\uparrow \mathbf{p}]$. Recall that a point is nothing but a dyadic box all of whose components are unit dyadic intervals.

The principal filter $\mathcal{P}_d^k[\uparrow \mathbf{p}]$ is isomorphic to the kth Cartesian power of a chain, each chain is of length d. It is well-known from poset theory that any Cartesian power of a chain is graded, rank symmetric, rank unimodal, and satisfies the *Sperner property* [19, 33]. In particular, the middle-rank set forms a maximum antichain.

For $0 \leqslant m \leqslant kd$, the mth-rank of the chain product $\mathcal{P}_d^k[\uparrow \mathbf{p}]$ has size equal to the number of integral solutions to the following equation

$$x_1 + \cdots + x_k = m, \;\; 0 \leqslant x_i \leqslant d, i \in [k].$$

By inclusion-exclusion, the number of such solutions is

$$\sum_{i=0}^{k} (-1)^i \binom{k}{i} \binom{m+k-1-id-i}{k-1}.$$

We should pick $m = kd/2$ because of rank-unimodality. The sum does not have a closed-form solution, even with the help of hypergeometric series. But the first term is bounded by $\binom{kd/2+k-1}{k-1}$. (Of course, one can also take the much simpler but weaker bound of d^k.) This is an upperbound on the number of maximal dyadic gap-boxes which contain a given probe point. The number is exponential in the query size, and up to poly-log in the input size, which we hide in the $\tilde{O}$-notation for brevity.[10] □

Remark B.13. For completeness, let us also discuss the case when the domain sizes are not uniform. If $|\mathbf{D}(A_i)| = D_i \leqslant 2^{d_i}$, then the set of all dyadic boxes is the product $\mathcal{P} = \mathcal{P}_{d_1} \times \cdots \times \mathcal{P}_{d_k}$. Then, the principal filter of $\mathcal{P}$ at $\mathbf{p}$ is a product of chains:

$$\mathcal{P}[\uparrow \mathbf{p}] = \mathcal{P}_{d_1}[\uparrow \mathbf{p}] \times \cdots \times \mathcal{P}_{d_k}[\uparrow \mathbf{p}].$$

The Cartesian product of k chains also satisfies the Sperner property; and, it is graded, rank-symmetric and rank-unimodal. The minimum rank is 0, and the maximum rank is $\sum_{i=1}^{k} d_i$. For any m such that

$$0 \leqslant m \leqslant d_1 + d_2 + \cdots + d_k,$$

the mth-rank of the poset $\mathcal{P}[\uparrow \mathbf{p}]$ has size equal to the number of integral solutions to the following equation

$$x_1 + \cdots + x_k = m, \;\; 0 \leqslant x_i \leqslant d_i, i \in [k].$$

By inclusion-exclusion, the number of such solutions is

$$\sum_{i=0}^{k} (-1)^i \sum_{S \in \binom{[k]}{i}} \binom{m+k-1-\sum_{j \in S} d_j - i}{k-1}.$$

This is a nasty sum. But it can always be upperbounded by the first term (or the first odd number of terms):

$$\binom{m+k-1}{k-1}.$$

To get the largest-sized rank, we pick

$$m = \left\lfloor \frac{d_1 + \cdots + d_k}{2} \right\rfloor.$$

[10] In fact, most of our algorithms has only a log-dependence on the input size instead of poly-log; but we will not make this distinction precise for the sake of clarity.

Proposition B.14. *Every (not necessarily dyadic) box in n dimensions can be decomposed into a disjoint union of at most* $(2d)^n = \tilde{O}(1)$ *dyadic boxes.*

Proof. It is sufficient to show that every closed interval (i.e. a 1-dimensional box) can be written as a disjoint union of at most $2d$ dyadic segments. Consider the complete binary tree where each node is labeled with the 01 path from the root down to the node; a left branch is labeled 0 and the right is labeled 1. Then, the set of all leaves represent the domain $\mathbf{D}(A)$ for an attribute A. Each closed interval is a set of consecutive leaves. Each node in the tree corresponds to a dyadic segment. Now, when we have a set of consecutive nodes at some level, we can merge every two sibling nodes into their parent, which is a dyadic segment containing the two siblings. Continue this merging process until no more merging is possible, then we get a collection of dyadic segments covering the original set of leaves. There are at most two nodes left per depth of this tree, because if there were 3 then some 2 of them can be merged. □

The following proposition follows immediately.

Proposition B.15. *For every box certificate, there is a dyadic box certificate of size at most a factor of* $(2d)^n = \tilde{O}(1)$ *larger.*

C Main Algorithm, Data Structure, and Analytical Idea

C.1 Dyadic data structure and dyadic resolution

There are two key operations that Tetris performs repeatedly: (1) querying for all boxes containing a given box, and (2) "resolving" two boxes to get a new box. We explain how a dyadic data structure can be used to perform those operations efficiently.

Data structure for storing dyadic boxes The operations in Line 1 of Algorithm 1, and Line 4 of Algorithm 2 are essentially the same operation: given a box $\mathbf{b}$, return the set of all dyadic boxes containing $\mathbf{b}$ from a storage of dyadic boxes. This operation can easily be supported in $\tilde{O}(1)$ time, as shown by Proposition B.12.

To store a collection of dyadic boxes, there are many options (including a hash table). We briefly describe here a natural implementation using a multi-level dyadic tree.

A *dyadic tree* is a binary tree storing dyadic segments, i.e. binary strings of length at most d. At each node, the left branch corresponds to bit 0 and the right corresponds to bit 1. Each time we insert a new dyadic segment x into the tree, we follow the bits of x down the tree. If we end up with a node already created in the tree, then we mark that node as a storage node. If we end up at a leaf while there are still bits left in x, then we create new nodes according to the bits of x that still need to be visited.

A dyadic tree can be used to store a set of dyadic segments. Recall that each dyadic box is an n-tuple of dyadic segments, i.e. an n-tuple of strings of the form $\mathbf{b} = \langle x_1, x_2, \ldots, x_n \rangle$, where $x_i \in \{0,1\}^{\leq d}$ for all $i \in [n]$. Fix some global attribute order $(A_1, A_2, \ldots, A_n)$. In a *multi-level dyadic tree*, we store a dyadic box $\langle x_1, \ldots, x_n \rangle$ by using a dyadic tree to store x_1, then at the storage node of x_1, there is a pointer to the root of a second dyadic tree storing x_2, etc. If there are two dyadic boxes sharing the same dyadic segment x_1, then their second dyadic segments are stored in the same tree, naturally. Figure 13 illustrates this simple data structure.

From Proposition B.12, the number of dyadic boxes containing $\mathbf{b}$ is at most $d^n = \tilde{O}(1)$. It is straightforward to see that the multilevel dyadic tree can return the set of all boxes containing a given dyadic box in $\tilde{O}(1)$ time.

Resolution Note that by the time Resolve is called in Line 18 of Algorithm 1, we know none of $\mathbf{w}_1$ and $\mathbf{w}_2$ covers $\mathbf{b}$. There are a lot of boxes we can infer from $\mathbf{w}_1$ and $\mathbf{w}_2$ if those two boxes are general dyadic boxes which can overlap in peculiar ways. However, Tetris forces $\mathbf{w}_1$ and $\mathbf{w}_2$ to be somewhat special.

Our Resolve routine uses ordered geometric resolution (recall Definition 4.2), and clearly $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ covers $\mathbf{b}$ as desired. We next prove the claimed property of $\mathbf{w}_1$ and $\mathbf{w}_2$.

Lemma C.1 (All resolutions are ordered)**.** *If the initial call to* TetrisSkeleton *is with the universal box* $\mathbf{b} = \langle \lambda, \ldots, \lambda \rangle$*, then the following hold:*

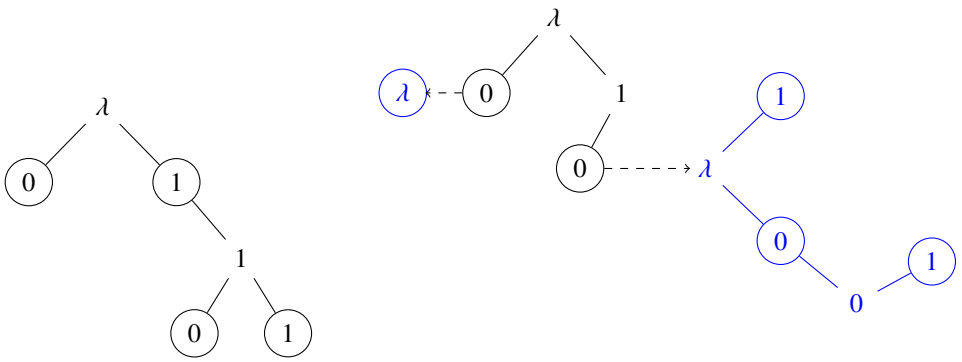

(a) Dyadic tree storing four dyadic segments 0, 1, 110, and 111

(b) Multilevel dyadic tree storing $\langle 0, \lambda\rangle$, $\langle 10, 1\rangle$, $\langle 10, 0\rangle$, $\langle 10, 001\rangle$

Figure 13: A simple data structure for storing dyadic boxes

(i) At any point in time, the box $\mathbf{b}$ *to be split must have the form*

$$\mathbf{b} \quad = \quad \langle x_1 \ , \ \ldots \ , \ x_{\ell-1} \ , \ x_\ell \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle \tag{7}$$

where $\ell \in [n]$ *and* $x_1, \ldots, x_{\ell-1}$ *are strings of length* d *(maximal-length strings), and* $|x_\ell| < d$.

(ii) every time we are calling Resolve $(\mathbf{w}_1, \mathbf{w}_2)$*, the two witnesses have the format shown in* (1) *and* (2).

In particular, it is sufficient for Resolve *to apply ordered resolution.*

Proof. We prove (i) by induction. The universal box certainly has the claimed format. When we split a box $\mathbf{b}$ (in line 6 of Algorithm 1) that has the format (7), $\mathbf{b}_1$ and $\mathbf{b}_2$ will be:

$$\begin{aligned} \mathbf{b}_1 \quad &= \quad \langle x_1 \ , \ \ldots \ , \ x_{\ell-1} \ , \ x_\ell 0 \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle \\ \mathbf{b}_2 \quad &= \quad \langle x_1 \ , \ \ldots \ , \ x_{\ell-1} \ , \ x_\ell 1 \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle \end{aligned}$$

This completes the proof of the invariant. To see (ii), note that when we call Resolve$(\mathbf{w}_1, \mathbf{w}_2)$, we know $\mathbf{w}_1$ covers $\mathbf{b}_1$ but not $\mathbf{b}$ (due to line 11 in Algorithm 1), and $\mathbf{w}_2$ covers $\mathbf{b}_2$ but not $\mathbf{b}$ (due to line 16 in Algorithm 1). It follows that $\mathbf{w}_1$ and $\mathbf{w}_2$ must be of the forms (1) and (2), respectively. □

Runtime is linear in the number of resolutions We first prove Lemma 4.3.

Remember that whenever Algorithm 1 does a resolution, it adds the result back to $\mathcal{A}$ (in Line 19). Therefore, M can be thought of as the total increase in the size of $\mathcal{A}$ since the initial call to TetrisSkeleton.

Proof of Lemma 4.3. We begin with the simplifying assumption that Tetris calls TetrisSkeleton only once. In this case if M is the number of resolutions made, we will prove that TetrisSkeleton (and hence Tetris) runs in time $\tilde{O}(M+1)$. Later on in the proof, we will see how to get rid of this assumption.

TetrisSkeleton can be thought of as a depth-first traversal of a binary tree whose nodes are dyadic boxes. Calling TetrisSkeleton ($\mathbf{b}$) corresponds to visiting node $\mathbf{b}$ in this binary tree. As soon as TetrisSkeleton visits a node $\mathbf{b}$, it checks whether $\mathbf{b}$ is covered by some dyadic box $\mathbf{a}$ in $\mathcal{A}$ and whether $\mathbf{b}$ is a point (lines 1 and 3 of Algorithm 1): If $\mathbf{b}$ is either covered or a point, then TetrisSkeleton backtracks directly without visiting any children of $\mathbf{b}$, in which case $\mathbf{b}$ is a *leaf node* of the visited binary tree. If $\mathbf{b}$ is neither covered nor a point, then TetrisSkeleton recursively visits its children $\mathbf{b}_1$ and $\mathbf{b}_2$, in which case $\mathbf{b}$ is an *internal node*. TetrisSkeleton spends $\tilde{O}(1)$ at each node it visits. In any full binary tree, the number of leaf nodes is equal to one plus the number of internal nodes. Therefore, to bound the runtime of TetrisSkeleton, we only need to bound the number of internal nodes it visits (i.e. the number of recursive calls of TetrisSkeleton in which the execution reaches line 6).

Moreover, whenever TetrisSkeleton returns from visiting a leaf with a FALSE (in line 4), then (thanks to lines 9 and 14) it will keep backtracking while passing this FALSE upwards the tree until it reaches the root. Because the depth of any node in the tree is $\tilde{O}(1)$, the total number of internal nodes that return FALSE is $\tilde{O}(1)$. (This term contributes to the $+1$ term in our final bound of $\tilde{O}(M+1)$.) Therefore, to bound the runtime of TetrisSkeleton, we only need to bound the number of internal nodes that return TRUE (i.e. the number of recursive calls that return in lines 12, 17 or 20). We will be referring to those nodes as TRUE *internal nodes*.

While TetrisSkeleton is traversing the tree, new boxes are being created by resolution and added to $\mathcal{A}$ (in lines 18 and 19). We will be referring to those boxes as *resolution boxes*. We will show that the total number of TRUE internal nodes is within a $\tilde{O}(1)$ factor from the total number of resolution boxes. And to do that, we will establish a mapping between TRUE internal nodes and resolution boxes such that this mapping satisfies the following two conditions:

- No resolution box is mapped to more than $\tilde{O}(1)$ TRUE internal nodes.
- Every TRUE internal node is mapped to at least one resolution box.

During the traversal, TetrisSkeleton makes sure not to visit any internal node $\mathbf{b}$ that is already covered by some box $\mathbf{a}$ in $\mathcal{A}$ (TetrisSkeleton might still visit $\mathbf{b}$ as a leaf node, but thanks to line 1, the execution will never make it to line 6 in order for $\mathbf{b}$ to become an internal node). However, when a new resolution box $\mathbf{w}$ is added to $\mathcal{A}$, the current internal node $\mathbf{b}$ that is being visited might already be covered by $\mathbf{w}$. If this happens, then (thanks to line 1) TetrisSkeleton will keep backtracking upwards the tree until it reaches a node that is not covered by $\mathbf{w}$, and it will never visit any internal node that is covered by $\mathbf{w}$ ever after. The depth of the tree is $\tilde{O}(1)$. Therefore, from the moment a new resolution box $\mathbf{w}$ is added to $\mathcal{A}$, TetrisSkeleton will traverse no more than $\tilde{O}(1)$ TRUE internal nodes that are covered by $\mathbf{w}$. We define the mapping between resolution boxes and TRUE internal nodes as follows: Every resolution box $\mathbf{w}$ is mapped to all TRUE internal nodes that are covered by $\mathbf{w}$ and that are traversed by TetrisSkeleton after $\mathbf{w}$ is added to $\mathcal{A}$. From this definition, we can see that no resolution box is mapped to more than $\tilde{O}(1)$ TRUE internal nodes.

Before TetrisSkeleton returns from visiting any internal node $\mathbf{b}$ with TRUE, $\mathbf{b}$ must be covered by some box in $\mathcal{A}$ (in lines 12, 17 and 20, $\mathbf{b}$ is covered by $\mathbf{w}_1, \mathbf{w}_2$ and $\mathbf{w}$, which all have been added to $\mathcal{A}$). Moreover, this box could not have existed in $\mathcal{A}$ from the very beginning because earlier in line 1 TetrisSkeleton could not find any box in $\mathcal{A}$ that covers $\mathbf{b}$. In other words, this box must be a resolution box that $\mathbf{b}$ can be mapped into. Therefore, every TRUE internal node is mapped to at least one resolution box.

Finally, we consider the case when Tetris calls TetrisSkeleton more than once. The proof of Lemma E.1 suggests that the total number of iterations of Tetris[11] (and hence the total number of calls to TetrisSkeleton) is bounded by $(|\mathcal{C}_\square| + Z)$. Note that the last call to TetrisSkeleton returns $(\text{TRUE}, \langle\lambda, \ldots, \lambda\rangle)$, indicating that $\langle\lambda, \ldots, \lambda\rangle$ is covered. At that time, Tetris would have implicitly built a resolution proof of $\langle\lambda, \ldots, \lambda\rangle$ using the Z output unit boxes along with at least $|\mathcal{C}_\square|$ input boxes. Hence, $|\mathcal{C}_\square| + Z = \tilde{O}(M)$.

□

The following definition comes up naturally in the runtime analysis of Tetris in Section 4.3 and later. It is just a generalization of the concept of prefix of a string.

Definition C.2 (Prefix of a box)**.** Given any box $\mathbf{b}$:

$$\mathbf{b} \quad = \quad \langle x_1 \ , \ \ldots \ , \ x_n \rangle$$

(where $x_1, \ldots, x_n$ are strings of arbitrary length), we define a *prefix* of $\mathbf{b}$ to be any box $\mathbf{b}'$ that has the form:

$$\mathbf{b}' \quad = \quad \langle x_1 \ , \ \ldots \ , \ x_{l-1} \ , \ \text{prefix}\,(x_l) \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle$$

where $l \in [n]$.

Notice that $\mathbf{b}'$ contains $\mathbf{b}$. It results from $\mathbf{b}$ by removing bits from the end. If we ignore the commas and consider all strings of each box to be a single string, then $\mathbf{b}'$ is indeed a prefix of $\mathbf{b}$.

[11] Even though Lemma E.1 is referring to Tetris-Reloaded, it can be easily verified that the bound of $(|\mathcal{C}_\square| + Z)$ on the number of iterations applies to Tetris in general.

In the context of joins (and BCP), when we resolve two boxes $\mathbf{w}_1$ and $\mathbf{w}_2$, each one of them could be either an input gap box or an output box (the latter is the box $\mathbf{b}$ added due to lines 7 and 8 in Algorithm 2) or a result of a previous resolution.

Definition C.3 (gap box resolution)**.** A *gap box resolution* is any resolution where each one of the two boxes to be resolved is either an input gap box or a result of a previous gap box resolution (i.e., a gap box resolution does not involve any output box, neither directly nor indirectly). The result of a gap box resolution is called a *gap box resolvent*, as defined in Section 4.4.

Definition C.4 (output box resolution)**.** An *output box resolution* is any resolution that is not a gap box resolution (i.e., at least one of the two boxes to be resolved is either an output box or a result of a previous output box resolution). The result of an output box resolution is called an *output resolvent*, as defined in Section 4.4.

The following theorem can be thought of as a template for almost all runtime bounds that are presented in this paper. We apply this theorem by bounding the number of gap box resolutions that Tetris performs depending on the input query's structural information.

Theorem C.5 (The gap box resolution bound)**.** *An invocation of* Tetris *runs in time $\tilde{O}(X + Z)$, where X is the total number of gap box resolutions that have been performed, and Z is the total number of output tuples of the* BCP.

Proof. We will use Lemma 4.3, and bound the runtime by bounding the number of resolutions. The number of gap box resolutions is X. Next, we bound the number of output box resolutions.

We will prove the following claim: For every witness $\mathbf{w}$ that is either an output box or an output resolvent, there is some output box $\mathbf{z}$ such that $\mathbf{w}$ is a prefix box of $\mathbf{z}$. Because every box has $\tilde{O}(1)$ possible prefix boxes, proving the claim will bound the number of output box resolutions by $\tilde{O}(Z)$, as desired.

This claim can be proved by induction. The base case clearly holds. For the inductive step, suppose that $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ where the claim holds for at least one of $\mathbf{w}_1$ or $\mathbf{w}_2$. Without loss of generality, let the claim hold for $\mathbf{w}_1$. $\mathbf{w}_1$ must be a prefix box of some output box $\mathbf{z}$. $\mathbf{z}$ is a unit box (i.e. a point) and hence has the form

$$\mathbf{z} \quad = \quad \langle t_1 \ , \ \ldots \ , \ t_n \rangle$$

where $t_1, \ldots, t_n$ are strings of maximal-length d. In Algorithm 2, the calls to TetrisSkeleton are with the universal box $\mathbf{b} = \langle \lambda, \ldots, \lambda \rangle$. According to Lemma C.1, all performed resolutions are going to be ordered. In particular, $\mathbf{w}_1$ and $\mathbf{w}_2$ must have the form shown in (1) and (2). Since $\mathbf{w}_1$ is a prefix box of $\mathbf{z}$, $\mathbf{w}_1$ and $\mathbf{w}_2$ will have the form

$$\begin{array}{rcccccccccccc} \mathbf{w}_1 & = & \langle & t_1 & , & \ldots & , & t_{l-1} & , & x_l 0 & , & \lambda & , & \ldots & , & \lambda & \rangle \\ \mathbf{w}_2 & = & \langle & \text{prefix}\,(t_1) & , & \ldots & , & \text{prefix}\,(t_{l-1}) & , & x_l 1 & , & \lambda & , & \ldots & , & \lambda & \rangle \end{array}$$

where $l \in [n]$. The output resolvent $\mathbf{w}$ will be

$$\mathbf{w} \quad = \quad \langle t_1 \ , \ \ldots \ , \ t_{l-1} \ , \ x_l \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle,$$

which is a prefix of $\mathbf{w}_1$ and hence $\mathbf{z}$. Therefore, the claim holds for $\mathbf{w}$. □

D Omitted details from Section 4.3

In this section, we show that by calling Tetris-Preloaded with an appropriate collection of gap boxes $\mathcal{B}(Q)$, we can achieve a runtime of the form $\tilde{O}(N + N^{\text{fhtw}} + Z)$, where N is the total input size, fhtw is the fractional hypertree width for the instance, and Z is the output size. When the query is α-acyclic, fhtw $= 1$, and thus Tetris-Preloaded runs in time linear in the input plus the output size. This is the celebrated Yannakakis result.

We will prove the above result by breaking it up into three steps. First, we show that Tetris-Preloaded achieves a run time of $\tilde{O}(N + \mathsf{AGM}(Q))$, where $\mathsf{AGM}(Q)$ is the best (i.e. tightest) AGM bound for the input instance. (See Section A.2.) Then, we show that Tetris-Preloaded runs in time $\tilde{O}(N + Z)$ if the query is α-acyclic. Finally, we prove the $\tilde{O}(N + N^{\text{fhtw}} + Z)$ result by using the first two results.

These results are conditioned on two key aspects of Tetris:

- The SAO used internally by TetrisSkeleton, which decides the dimension order in which the algorithm splits a target box **b**.
- The type of gap boxes that we chose to pre-load Tetris-Preloaded with (i.e., the set $\mathcal{B}$ in Algorithm 2.)

D.1 Tetris-Preloaded achieves AGM bound

Definition D.1 ($\hat{\mathcal{T}}(\mathcal{A})$: a multilevel dyadic tree pruned by a set of boxes $\mathcal{A}$)**.** Let $\mathcal{A}$ be a set of dyadic boxes. The multilevel dyadic tree that is pruned by $\mathcal{A}$, denoted by $\hat{\mathcal{T}}(\mathcal{A})$, is the set of all boxes **b** that have the format described by (7) and that satisfy the following condition:

- For every box $\mathbf{b}'$ that is a prefix box [12] of **b** such that $\mathbf{b}' \neq \mathbf{b}$ and for every box $\mathbf{a} \in \mathcal{A}$, we have $\mathbf{b}' \not\subseteq \mathbf{a}$ (i.e., $\mathbf{b}'$ is not covered by **a**).

Theorem D.2 (Tetris-Preloaded achieves AGM bound)**.** *Let Q be a join query, N the total number of input tuples, $\mathsf{AGM}(Q)$ the best AGM bound for this instance, and $\mathcal{B}(Q) = \bigcup_{R \in atoms(Q)} \mathcal{B}(R)$ the set of all input gap boxes. Then,* Tetris-Preloaded *runs in time $\tilde{O}(N + \mathsf{AGM}(Q))$ and $\left|\hat{\mathcal{T}}(\mathcal{B}(Q))\right|$ is bounded by $\tilde{O}(\mathsf{AGM}(Q))$, provided that the following conditions are met:*

- *The SAO σ used by the algorithm is arbitrary, but fixed.*
- *For each input relation $R \in atoms(Q)$, $\mathcal{B}(R)$ contains only σ-consistent gap boxes. (This condition can easily be met if the search tree structure for R is consistent with σ.)*

Proof. The term $\tilde{O}(N)$ is the time needed to initialize $\mathcal{A}$ with all boxes of $\mathcal{B}$ (in line 1 of Algorithm 2). While explaining the proof, we will assume that the data structure that is being used to store boxes of $\mathcal{A}$ is a multilevel dyadic tree (Section C.1).

We will consider a slightly different version of TetrisSkeleton, denoted by TetrisSkeleton2. The difference is that whenever TetrisSkeleton2 encounters a unit box **b** that is not covered by any box in $\mathcal{A}$, instead of returning $(\textsc{false}, \mathbf{b})$ as TetrisSkeleton would do, TetrisSkeleton2 reports **b** as an output, adds **b** to $\mathcal{A}$, and returns $(\textsc{true}, \mathbf{b})$. In particular, instead of line 4 of TetrisSkeleton (Algorithm 1), TetrisSkeleton2 has the following three lines:

Report b as an output tuple
$\mathcal{A} \leftarrow \mathcal{A} \cup \{\mathbf{b}\}$
Return $(\textsc{true}, \mathbf{b})$

Unlike TetrisSkeleton which returns one output point per call, a single call to TetrisSkeleton2 reports all output points. First, we will explain the proof under the assumption that we are using TetrisSkeleton2. Later on, we will see how the proof holds for TetrisSkeleton as well.

We use the proof strategy from [51], implicit in [50]: by induction on the total number of attributes, then apply Hölder inequality to the inductive step. To bound the total runtime, it is sufficient to bound the number of boxes **b** that the algorithm considers. Each box **b** is nothing but a node in a multilevel dyadic tree, one level for each component of **b**. Hence, in the proof of this theorem, we will often speak of TetrisSkeleton2 "visiting" a node in the dyadic tree. "Visiting a node" means the current box **b** in TetrisSkeleton2 corresponds to that node.

The base case is when there is $n = 1$ attribute A and m unary relations, i.e. $|\text{atoms}(Q)| = m$. The join problem reduces to the problem of computing the intersection of m sets of strings of length d. The gaps are dyadic segments, i.e. strings of length at most d. When we insert gaps (from $\mathcal{B}(Q)$) into the internal dyadic data structure of TetrisSkeleton2, each gap corresponds to a node in the dyadic tree. We call nodes storing gaps *storage nodes*. All leaves of the dyadic tree are storage nodes (while internal nodes of the dyadic tree *could* be storage nodes). TetrisSkeleton2 starts from the root of the tree, visiting the left and right branches of each node recursively. If a storage node is hit, then the algorithm backtracks.

Tuples of any relation are unit boxes; we will be referring to them as *tuple boxes*. Now, consider the smallest relation R from the input. The gap boxes and tuple boxes of R correspond to storage nodes and depth-d leaves of the dyadic tree. (We do not store those depth-d leaves.) When TetrisSkeleton2 visits a storage node corresponding to a

[12] Check Definition C.2.

gap from R, it will backtrack. It might have backtracked earlier because it might have hit a storage node corresponding to a gap from a different relation. In either case, each gap box from R is visited at most once. If TetrisSkeleton2 hits a depth-d leaf corresponding to a tuple from R, it might or might not report this tuple as an output depending on whether it corresponds to a gap from a different relation or not. But either way, it will backtrack and never visit this leaf again. If TetrisSkeleton2 does not hit this leaf in the first place, it must have backtracked earlier at a parent node of this leaf. There are at most d such parent nodes. In summary, the total runtime of TetrisSkeleton2 is at most linear in the number of tuples from R plus the number of gap boxes from R. Thanks to Proposition B.12, the total runtime is bounded by $\tilde{O}(|R|)$.

Next, let $\{x_S\}_{S\in\text{atoms}(Q)}$ be a fractional edge cover of the hypergraph of this query, then $\sum_{S\in\text{atoms}(Q)} x_S \geqslant 1$ and $x_S \geqslant 0$ for all $S \in \text{atoms}(Q)$. The overall runtime is $\tilde{O}$ of

$$\min_{S\in\text{atoms}(Q)} |S| = |R| \leqslant |R|^{\sum_{S\in\text{atoms}(Q)} x_S} \leqslant \prod_{S\in\text{atoms}(Q)} |S|^{x_S}.$$

The right-hand side is the AGM bound that corresponds to the chosen fractional edge cover. (See Equation (6).) The base case is thus proved.

Next, consider the general case when $n > 1$. Without loss of generality, assume $\sigma = (A_1, \ldots, A_n)$. Recall that a σ-consistent gap box $\mathbf{b} = \langle x_1, \ldots, x_n\rangle$ has the following property: there is one position $i \in [n]$ for which $|x_j| = 0$ (i.e. $x_j = \lambda$) for every $j > i$, and $|x_j| \in \{0, d\}$ for every $j < i$. (Definition 3.10.)

We bound the number of boxes $\mathbf{b} = \langle x_1, \ldots, x_n\rangle$ visited by TetrisSkeleton2 by considering two types of $\mathbf{b}$: (type 1) the boxes $\mathbf{b}$ for which $x_2 = \lambda$, and (type 2) the boxes $\mathbf{b}$ for which $x_2 \neq \lambda$. Note that by Lemma C.1, if $x_2 = \lambda$ then $x_i = \lambda$ for all $i > 2$.

To bound the number of type-1 boxes TetrisSkeleton2 visits, we only have to note that TetrisSkeleton2 visits these nodes exactly as if it was computing $\bowtie_{R:A_1\in\text{vars}(R)} \pi_{A_1}(R)$. To see this, consider any value $t_1 \notin \bowtie_{R:A_1\in\text{vars}(R)} \pi_{A_1}(R)$. Then, $t_1 \notin \pi_{A_1}(R)$ for some R. In that case, out of the σ-consistent gap boxes from R, there must be a gap box of the form $\mathbf{a} = \langle a_1, \lambda, \ldots, \lambda\rangle$ where a_1 is a prefix of t_1. And thus, when the algorithm hits the storage node corresponding to a_1 at the first-level dyadic tree, it will have to backtrack. Hence, as shown in the base case, the total number of type-1 boxes TetrisSkeleton2 visits is bounded by (within a $\tilde{O}(1)$ factor)

$$\min_{R:A_1\in\text{vars}(R)} |R| \leqslant \prod_{R:A_1\in\text{vars}(R)} |R|^{x_R} \leqslant \text{AGM}(Q).$$

To bound the number of type-2 boxes TetrisSkeleton2 visits, we note that each such box $\mathbf{b} = \langle x_1, x_2, \ldots, x_n\rangle$ must have $|x_1| = d$ (Lemma C.1) and $x_1 \in \bowtie_{R:A_1\in\text{vars}(R)} \pi_{A_1}(R)$. For every value $x_1 \in \bowtie_{R:A_1\in\text{vars}(R)} \pi_{A_1}(R)$, TetrisSkeleton2 will be visiting type-2 boxes $\mathbf{b}$ that start with this x_1 (i.e. $\mathbf{b} = \langle x_1, x_2, \ldots, x_n\rangle$) exactly as if it was computing the join

$$\left(\bowtie_{R:A_1\in\text{vars}(R)} \pi_{A_2,\ldots,A_n}\sigma_{A_1=x_1}R\right) \bowtie \left(\bowtie_{R:A_1\notin\text{vars}(R)} R\right).$$

The reason is that relations that contain A_1 provide gap boxes of the form $\mathbf{a} = \langle x_1, a_2, \ldots, a_n\rangle$, while other relations provide gap boxes of the form $\mathbf{a} = \langle \lambda, a_2, \ldots, a_n\rangle$. All those gap boxes are pre-loaded in $\mathcal{A}$. Once TetrisSkeleton2 visits a box $\mathbf{b}$ that is covered by any box $\mathbf{a}$ in $\mathcal{A}$, it will backtrack (thanks to Line 1 of Algorithm 1).

To sum up, TetrisSkeleton2 computes the intersection of the projections of all relations on A_1. And for each value x_1 in this intersection, TetrisSkeleton2 applies itself recursively on the selection of x_1 from all relations. This behavior is identical to that of the generic join algorithm (Algorithm 3 in [51]). The inductive step follows from the query decomposition lemma (i.e. Lemma 3.1 from [51] with $|I| = 1$). Hence, we have proved that the number of boxes $\mathbf{b}$ visited by TetrisSkeleton2 is $\tilde{O}(\text{AGM}(Q))$.

Every box $\mathbf{b}$ TetrisSkeleton2 visits must belong to $\hat{\mathcal{T}}(\mathcal{B}(Q))$, but the converse is not necessarily true. However, the above proof not only bounds the number of boxes $\mathbf{b}$ visited by TetrisSkeleton2, but also the superset $\hat{\mathcal{T}}(\mathcal{B}(Q))$.

Finally, we remove our assumption about TetrisSkeleton2. In particular, we show that using TetrisSkeleton instead is not going to slow down the outer algorithm Tetris-Preloaded by more a $\tilde{O}(1)$ factor. Before TetrisSkeleton returns from visiting a node $\mathbf{b}$ with TRUE, it adds a box $\mathbf{w}$ containing $\mathbf{b}$ to $\mathcal{A}$. This way, TetrisSkeleton makes sure it never visits any children of $\mathbf{b}$ ever after. Suppose that TetrisSkeleton encounters an output point $\mathbf{o}$. Then, it returns directly to the root with FALSE. The next call of TetrisSkeleton will be heading straight towards $\mathbf{o}$, thanks to the boxes $\mathbf{w}$ that

were added to $\mathcal{A}$ by the previous call. Hence, within a $\tilde{O}(1)$ time, the next call will be resuming the work of the previous call, as if the two were a single continues call. As a result, the multiple calls to TetrisSkeleton are equivalent to a single call to TetrisSkeleton2. □

If we prevent Tetris from caching resolution results (i.e. if we drop line 19 of Algorithm 1), then Tetris will be performing what is known as *tree* resolution: Every box that results from a resolution will become an input to maximally one resolution. Hence, the resolution proof generated by Tetris will have a tree structure. Because $\hat{\mathcal{T}}(\mathcal{B}(Q)) = \tilde{O}(\mathsf{AGM}(Q))$, Tetris-Preloaded without caching can still achieve the bound of $\tilde{O}(N + \mathsf{AGM}(Q))$. [13]

Corollary D.3 (Tree Ordered Geometric Resolution achieves AGM bound)**.** *Let Q be a join query, N the total number of input tuples, $\mathsf{AGM}(Q)$ be the best AGM bound for this instance, and let the conditions of Theorem D.2 be met. Then, Q can be solved using $\tilde{O}(\mathsf{AGM}(Q))$ tree ordered geometric resolutions.*

D.2 Tetris-Preloaded matches Yannakakis algorithm on α-acyclic queries

To show the main result of this section, we need a simple auxiliary lemma. Recall that a GYO-elimination order [1] is obtained by repeating the following two operations on the hypergraph of an α-acyclic query: (a) remove a vertex that belongs to only one hyperedge (i.e. a *private* vertex), (b) remove a hyperedge which is a subset of another hyperedge. As mentioned earlier in Definition A.3, a hypergraph is α-acyclic if and only if repeating the above two operations arbitrarily results in an empty hypergraph. The order in which vertices are removed is called a *GYO elimination order*. Notice that there can be many GYO elimination orders for a given α-acyclic query. For example, if the query has two relations $R(A_1, \ldots, A_n)$ and $S(A_1, \ldots, A_n)$, then there are $n!$ GYO elimination orders.

Lemma D.4. *Let Q be an α-acyclic query. Let $\sigma = (A_1, \ldots, A_n)$ be the reverse of any GYO elimination order for Q. For any $k \in [n]$, define $\mathcal{R}_k = \{R \in atoms(Q) \mid A_k \in \text{vars}(R)\}$. Then, there exists a relation $R \in \mathcal{R}_k$ satisfying the following property:*

$$\forall S \in \mathcal{R}_k \qquad [\text{vars}(S) \cap \{A_1, \ldots, A_k\} \quad \subseteq \quad \text{vars}(R) \cap \{A_1, \ldots, A_k\}] \tag{8}$$

Proof. When A_n is first removed by GYO-elimination, it must be contained in only one hyperedge e of the hypergraph of Q. All other hyperedges containing A_n must have been removed earlier. Therefore, all those hyperedges must have been contained in other hyperedges that contain A_n. Hence, the relation R with $\text{vars}(R) = e$ satisfies property (8). After A_n is removed, the residual graph is α-acyclic. Induction completes the proof. □

Definition D.5 (Bottom relation and $\mathsf{support}(A_k)$)**.** Fix a GAO that is any reversed GYO elimination order. We will refer to a relation $R \in \mathcal{R}_k$ satisfying condition (8) as a *bottom relation for A_k* with respect to the GAO. The set $\text{vars}(R) \cap \{A_1 \ldots, A_k\}$ is called the *support for A_k*, denoted by $\mathsf{support}(A_k)$. Another way to state condition (8) is to say that the union of all the sets $\text{vars}(S) \cap \{A_1, \ldots, A_k\}$ is a subset of $\mathsf{support}(A_k)$.

Note that there can be multiple bottom relations for A_k. For example, in the query $Q = R(A, B, C) \bowtie S(A, B, C) \bowtie T(A, B, C)$, all the relations are bottom relations for A, for B, and for C.

Definition D.6 (Bottom boxes for A_k)**.** Fix the GAO σ to be any reversed GYO elimination order, and consider only σ-consistent gap boxes from all relations. Let $\mathcal{N}_k$ denote the set of all input gap boxes *and* input tuple boxes $\mathbf{b}$ whose supports are precisely $\mathsf{support}(A_k)$.

Obviously, the boxes from $\mathcal{N}_k$ can only come from the bottom relations for A_k; and only the gap boxes (and not the tuple boxes) from the bottom relations are part of the input to Tetris-Preloaded. The following proposition is straightforward.

[13] For that to hold, the step of reporting output points has to be moved from the outer Tetris-Preloaded into TetrisSkeleton, the same way we did in the proof of Theorem D.2 when we defined TetrisSkeleton2.

Proposition D.7. *Fix the GAO σ to be any reversed GYO elimination order, and consider only σ-consistent gap boxes from all relations. For any $k \in [n]$, the number of bottom boxes for A_k is linear (within a $\tilde{O}(1)$ factor) in the total number of tuples from all the bottom relations for A_k. In particular,*

$$\sum_{k=1}^{n} |\mathcal{N}_k| = \tilde{O}(N).$$

Theorem D.8 (Tetris-Preloaded matches Yannakakis algorithm)**.** *Let Q be an α-acyclic join query, N the total number of input tuples, and Z total number of output tuples. Then,* Tetris-Preloaded *runs in time $\tilde{O}(N + Z)$, provided that the following conditions are met:*

- *The SAO σ used by the algorithm is the reverse of some GYO elimination order.*[14]
- *For each $R \in atoms(Q)$, the boxes in $\mathcal{B}(R)$ are σ-consistent.*

Proof. WLOG, assume $\sigma = (A_1, \ldots, A_n)$. We prove this theorem by applying Theorem C.5 and bounding the number of gap box resolutions to be $\tilde{O}(N)$. Thanks to Proposition D.7, it is sufficient to show the following claim:

Claim 1. For any $k \in [n]$, the number of gap box resolutions on A_k is $\tilde{O}(|\mathcal{N}_k|)$.

We first show that the claim holds for $k = n$, i.e. the number of gap box resolutions on A_n is $\tilde{O}(|\mathcal{N}_n|)$. Our plan is to show that every gap box resolution on A_n produces a box $\mathbf{w}$ which *covers* some box $\mathbf{a} \in \mathcal{N}_n$. From Proposition B.12, the number of such boxes $\mathbf{w}$ is thus bounded by $\tilde{O}(|\mathcal{N}_n|)$.

Resolution occurs only in line 18 of Algorithm 1. For completeness, we repeat some of the arguments leading to Lemma C.1. For the witnesses $\mathbf{w}_1$ and $\mathbf{w}_2$ to resolve on the last attribute A_n, their n-th components must be $x_n 0$ and $x_n 1$ respectively for some string x_n. We only need to resolve $\mathbf{w}_1$ and $\mathbf{w}_2$ when each one of them covers one half of $\mathbf{b}$ but neither one of them covers $\mathbf{b}$ as a whole. This only happens when $\mathbf{b}, \mathbf{b}_1, \mathbf{b}_2$ have the forms

$$\begin{aligned} \mathbf{b} &= \langle x_1 \ , \ \ldots \ , \ x_{n-1} \ , \ x_n \rangle \\ \mathbf{b}_1 &= \langle x_1 \ , \ \ldots \ , \ x_{n-1} \ , \ x_n 0 \rangle \\ \mathbf{b}_2 &= \langle x_1 \ , \ \ldots \ , \ x_{n-1} \ , \ x_n 1 \rangle. \end{aligned}$$

where $|x_i| = d$ for all $i < n$. When this is the case, then $\mathbf{w}_1, \mathbf{w}_2$ and $\mathbf{w}$ must have the forms

$$\begin{aligned} \mathbf{w}_1 &= \langle \text{prefix}\,(x_1) \ , \ \ldots \ , \ \text{prefix}\,(x_{n-1}) \ , \ x_n 0 \rangle \\ \mathbf{w}_2 &= \langle \text{prefix}\,(x_1) \ , \ \ldots \ , \ \text{prefix}\,(x_{n-1}) \ , \ x_n 1 \rangle \\ \mathbf{w} &= \langle \text{prefix}\,(x_1) \ , \ \ldots \ , \ \text{prefix}\,(x_{n-1}) \ , \ x_n \rangle. \end{aligned}$$

where one prefix in $\mathbf{w}_1$ can have a different length than the corresponding prefix in $\mathbf{w}_2$, and $\mathbf{w}$ will inherit the longer prefix between the two (i.e. the one that corresponds to the intersection of the two dyadic segments).

Positive witnesses [15] can be obtained in three ways: either from input gap boxes, or from output tuples, or by the (ordered) resolution of other witnesses (The first two cases can occur in Line 2 while the third one occurs in Line 18 of Algorithm 1). Since we are only considering gap box resolutions, the witnesses $\mathbf{w}_1$ and $\mathbf{w}_2$ do not contain any output tuple; in particular, each of them is either an input gap box, or is a result of previous resolutions on input gap boxes.

To prove Claim 1 (for $k = n$), we will show that $\mathbf{w}$ covers some box in the set $\mathcal{N}_n$. Note that, by definition, $\mathcal{N}_n$ contains either gap boxes or tuple boxes from the bottom relations of A_n. To show that $\mathbf{w}$ covers some box in $\mathcal{N}_n$, we show by induction that every witness involved in a resolution on A_n must cover some box in $\mathcal{N}_n$. For the base case, suppose $\mathbf{w}_1$ and $\mathbf{w}_2$ are input gap boxes. We want to show that each one of them covers a box in $\mathcal{N}_n$. Note that because they are input gap boxes, $\mathsf{support}(\mathbf{w}_1) \subseteq \mathsf{support}(A_n)$ and $\mathsf{support}(\mathbf{w}_2) \subseteq \mathsf{support}(A_n)$.

Let $\overline{\mathbf{b}}$ denote the projection [16] of $\mathbf{b}$ onto the support of A_n, i.e. $\overline{\mathbf{b}} = \pi_{\mathsf{support}(A_n)}(\mathbf{b})$. Since no input gap boxes cover $\mathbf{b}$ (otherwise we would not have called Split-First-Thick-Dimension($\mathbf{b}$)), no input gap boxes cover $\overline{\mathbf{b}}$ either. In particular,

[14] i.e., the vertex that is eliminated first in GYO should be the last in the SAO σ.

[15] $\mathbf{w}$ is a positive witness iff it covers $\mathbf{b}$. Remember that in Tetris-Preloaded, input gap boxes and output tuples cover the whole space.

[16] See Definition E.2 for the notion of projection of a box onto some support.

each one of the boxes $\overline{\mathbf{b}}$, $\overline{\mathbf{b}_1} = \pi_{\mathsf{support}(A_n)}(\mathbf{b}_1)$ and $\overline{\mathbf{b}_2} = \pi_{\mathsf{support}(A_n)}(\mathbf{b}_2)$ must be a prefix box of some box in $\mathcal{N}_n$. (See Definition C.2.) Because $\mathsf{support}(\mathbf{w}_1) \subseteq \mathsf{support}(A_n)$, $\mathbf{w}_1$ (which covers $\mathbf{b}_1$) covers $\overline{\mathbf{b}_1}$ as well. As a result, $\mathbf{w}_1$ covers some box in $\mathcal{N}_n$. And the same holds for $\mathbf{w}_2$.

Now, to complete the induction, we show that $\mathbf{w}$ (which is the resolution of $\mathbf{w}_1$ and $\mathbf{w}_2$) must cover some box in $\mathcal{N}_n$. Notice that

$$\mathsf{support}(\mathbf{w}) \subseteq \mathsf{support}(\mathbf{w}_1) \cup \mathsf{support}(\mathbf{w}_2) \subseteq \mathsf{support}(A_n)$$

Therefore, $\mathbf{w}$ (which covers $\mathbf{b}$) covers $\overline{\mathbf{b}}$ as well. But $\overline{\mathbf{b}}$ is a prefix of some box in $\mathcal{N}_n$. As a result, $\mathbf{w}$ covers some box in $\mathcal{N}_n$.

Next, we prove Claim 1 above for some $k < n$. The proof is almost the same as the $k = n$ case. However, we will need one additional claim, which can be proved by induction.

Claim 2. Every witness $\mathbf{w}$ that is involved in a gap box resolution on attribute A_k must satisfy $\mathsf{support}(\mathbf{w}) \subseteq \mathsf{support}(A_k)$.

Before proving Claim 2, let us see how it helps complete the proof of Claim 1. According to Lemma C.1, when we resolve $\mathbf{w}_1$ and $\mathbf{w}_2$ on A_k, all the attributes $A_{k+1}, \ldots, A_n$ will be λ's. Each one of the boxes $\mathbf{b}, \mathbf{b}_1, \mathbf{b}_2, \mathbf{w}_1, \mathbf{w}_2, \mathbf{w}$ will have $(n-k)$ trailing λs. For example,

$$\begin{aligned}
\mathbf{b} &= \langle x_1 \ , \ \ldots \ , \ x_{k-1} \ , \ x_k \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle \\
\mathbf{b}_1 &= \langle x_1 \ , \ \ldots \ , \ x_{k-1} \ , \ x_k 0 \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle \\
\mathbf{b}_2 &= \langle x_1 \ , \ \ldots \ , \ x_{k-1} \ , \ x_k 1 \ , \ \lambda \ , \ \ldots \ , \ \lambda \rangle.
\end{aligned}$$

Now, let $\overline{\mathbf{b}'} = \pi_{\mathsf{support}(A_k)}(\mathbf{b})$. Then no input gap boxes cover $\overline{\mathbf{b}'}$, otherwise we would not have split $\mathbf{b}$. In particular, each one of the boxes $\overline{\mathbf{b}'}$, $\overline{\mathbf{b}'_1} = \pi_{\mathsf{support}(A_k)}(\mathbf{b}_1)$ and $\overline{\mathbf{b}'_2} = \pi_{\mathsf{support}(A_k)}(\mathbf{b}_2)$ must be a prefix box of some box in $\mathcal{N}_k$. From Claim 2, $\mathsf{support}(\mathbf{w}_1) \subseteq \mathsf{support}(A_k)$ and $\mathsf{support}(\mathbf{w}_2) \subseteq \mathsf{support}(A_k)$. Consequently, both $\mathbf{w}_1$ and $\mathbf{w}_2$ cover boxes in $\mathcal{N}_k$. However,

$$\mathsf{support}(\mathbf{w}) \subseteq \mathsf{support}(\mathbf{w}_1) \cup \mathsf{support}(\mathbf{w}_2) \subseteq \mathsf{support}(A_k) \tag{9}$$

Consequently, $\mathbf{w}$ covers some box in $\mathcal{N}_k$ as well. (9) proves the inductive step of Claim 2 within the same value of k. For different values of k, Claim 2 can be proved inductively from $k = n$ down to 1. □

D.3 Tetris-Preloaded matches the fractional hypertree width bound

To obtain a runtime of $\tilde{O}(N + N^{\text{fhtw}} + Z)$, a typical strategy is as follows. We first compute a tree decomposition (Definition A.4) of the query for which the the maximum AGM bound over the bags is minimized. Then, we compute a set of intermediate relations, one for each bag, using an AGM-bound matching algorithm such as [50] or [71]. Finally, we run Yannakakis algorithm on the resulting "bag relations", because those relations form an α-acyclic query. Since each intermediate relation is of size at most N^{fhtw}, we have the claimed runtime.

The above strategy is essentially to "hide" the non α-acyclic parts inside bags. The strategy is somewhat unsatisfactory as we have to run the algorithm in two stages. We will show that Tetris-Preloaded can achieve the objective in "one shot."

We have already seen (in Theorems D.2 and D.8) that Tetris-Preloaded achieves AGM bound and is capable of playing the role of Yannakakis algorithm. Next, we will show that it is even more "polymorphic" than that: When applied on a join query given some tree decomposition, it will achieve AGM bound on each bag and simulate Yannakakis on the bags. In other words, we don't have to apply Tetris-Preloaded on each bag individually, and then one more time on all the bags together. The algorithm is "smart" enough to produce all this behavior when it is applied directly on the original query.

Given a bag G of some tree decomposition of a join query Q, we will use $\text{vars}(G)$ to denote the set of all attributes in G.

Theorem D.9 (Tetris-Preloaded achieves AGM and Yannakakis together)**.** *Let Q be a join query, N the total number of input tuples, Z the total number of output tuples, TD some tree decomposition of Q, and $\overline{\mathsf{AGM}_{TD}(Q)}$ the maximum over all TD bags of the best AGM bound for this bag. Then,* Tetris-Preloaded *runs in time $\tilde{O}(N + \overline{\mathsf{AGM}_{TD}(Q)} + Z)$, provided that the following conditions are met:*

- *The SAO σ used by the algorithm is the reverse of some GYO elimination order on the bags of TD.*[17]
- *For each relation $R \in atoms(Q)$, the boxes in $\mathcal{B}(R)$ are σ-consistent.*
- *For each bag G of TD, there exists a set of relations $atoms(G) \subseteq atoms(Q)$ that achieves the best* AGM *bound for G and that satisfies* $\text{vars}(R) \subseteq \text{vars}(G)$ *for all $R \in atoms(G)$.* [18]

Proof. WLOG, assume $\sigma = (A_1, \ldots, A_n)$. We will use the same outline of the proof of Theorem D.8. We start with bounding the total number of gap box resolutions on the last attribute A_n. Bags form an α-acyclic hypergraph. A_n is the first vertex/attribute that is eliminated in GYO. Consider all bags that contain A_n. At least one of them (let's call it G_n) must contain all the attributes that are contained in any of them (Lemma D.4). Hence, G_n is a *bottom bag* for A_n (in the sense of Definition D.5). Let AGM_n be the best AGM bound for G_n. We define a set $\mathcal{N}_n$ of size $\tilde{O}(\mathsf{AGM}_n)$. The plan is to show that every gap box resolution on A_n produces a box that covers some box in $\mathcal{N}_n$.

Let $\mathcal{B}(G_n)$ denote the set of input gap boxes from all relations in $\text{atoms}(G_n)$ projected onto $\text{vars}(G_n)$. The set $\mathcal{N}_n$ is chosen to be $\hat{\mathcal{T}}(\mathcal{B}(G_n))$. (See Definition D.1.) By Theorem D.2, $|\mathcal{N}_n| = \tilde{O}(\mathsf{AGM}_n)$.

While running Tetris-Preloaded on Q, consider the parameter $\mathbf{b}$ of TetrisSkeleton. Let $\overline{\mathbf{b}}$ be the projection of $\mathbf{b}$ onto $\text{vars}(G_n)$. It is not hard to see that $\overline{\mathbf{b}}$ must belong to $\mathcal{N}_n$.

To prove that every gap box resolution on A_n produces a box that covers some box in $\mathcal{N}_n$, we can prove two things: First, every input gap box that is used as a witness on A_n covers some box in $\mathcal{N}_n$ (and that box is $\overline{\mathbf{b}}$). Second, every resolution on A_n between two boxes that cover boxes in $\mathcal{N}_n$ produces another box that covers some box in $\mathcal{N}_n$ (and that box is also $\overline{\mathbf{b}}$).

Inductively, we can extend this to any attribute A_k for $k < n$. □

The following corollary is immediate.

Corollary D.10 (Tetris-Preloaded achieves fractional hypertree width bound)**.** *Let Q be a join query, N the total number of input tuples, fhtw be the fractional hypertree width of the query, and Z the total number of output tuples. Then,* Tetris-Preloaded *runs in time $\tilde{O}(N + N^{fhtw} + Z)$, provided that the following conditions are met:*

- *The SAO σ used by the algorithm is the reverse of some GYO elimination order on the bags of some tree decomposition TD whose fractional hypertree width is fhtw.*
- *For each relation $R \in atoms(Q)$, the box set $\mathcal{B}(R)$ is σ-consistent.*
- *For each bag G of TD, there exists a set of relations $atoms(G) \subseteq atoms(Q)$ that forms an optimal fractional edge cover of G and that satisfies* $\text{vars}(R) \subseteq \text{vars}(G)$ *for all $R \in atoms(G)$.* [19]

E Omitted details from Section 4.4

In this section, we consider join queries (and BCP instances) whose treewidth is w. We show that Tetris-Reloaded can be used to solve those queries in time $\tilde{O}(|\mathcal{C}_\square|^{w+1} + Z)$. Moreover, we show that Tetris-Reloaded runs in time $\tilde{O}(|\mathcal{C}_\square|^{w} + Z)$ in two special cases: when $w = 1$ and $w = n - 1$. However, when $1 < w < n - 1$, Theorem 5.3 shows a lowerbound of $\Omega(|\mathcal{C}_\square|^{w+1} + Z)$ on Ordered Geometric Resolution proof sizes, and hence a similar lower bound on the runtime of Tetris-Reloaded. Moreover, when $w = n - 1$, Theorem 5.4 shows a lower bound of $\Omega(|\mathcal{C}_\square|^{n-1} + Z)$.

To analyze the runtime of Tetris-Reloaded, we first argue that the number of accesses it makes to $\mathcal{B}$ is of the correct order:

Lemma E.1. *The total number of boxes that* Tetris-Reloaded *loads from $\mathcal{B}$ into $\mathcal{A}$ (in line 4) is $\tilde{O}(|\mathcal{C}_\square|)$.*

[17] The bags of any tree decomposition form an α-acyclic hypergraph. Hence, we can apply GYO elimination on this hypergraph by repeating those two steps in any order: (1) remove any bag that is contained in another bag, (2) remove any vertex/attribute that appears in only one bag.

[18] This condition can always be enforced as follows: If relation R is necessary to achieve AGM bound on bag G while $\text{vars}(R) \not\subseteq \text{vars}(G)$, then $\pi_{\text{vars}(G)}R$ has to be precomputed, added to $\text{atoms}(Q)$, and has to replace R in AGM bound of G.

[19] As before, precomputing $\pi_{\text{vars}(G)}R$ can always enforce this condition when needed.

Proof. In each iteration of Tetris-Reloaded, if $\mathbf{w}$ is an output point, then no box from $\mathcal{B}$ is loaded into $\mathcal{A}$. If $\mathbf{w}$ is not an output point, there must be at least one box from $C_\square$ that covers $\mathbf{w}$. Moreover, no box from $C_\square$ is loaded twice. Hence, the total number of iterations that load boxes from $\mathcal{B}$ is at most $|C_\square|$. The number of boxes that are loaded from $\mathcal{B}$ in each iteration (i.e. the size of $\mathcal{B}'$) is $\tilde{O}(1)$, thanks to Proposition B.12 in Appendix B.3, completing the proof. □

From Lemma 4.3, the runtime of Tetris is in the order of the number of resolutions $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ it performs. We call the box $\mathbf{w}$ a *resolvent* (box). It is not hard to see that a box cannot be a resolvent twice. Hence, it is sufficient to bound the total number of resolvents. We develop a simple technique for bounding the total number of resolvents Tetris encounters. This technique will be used many times to prove other runtime bounds in the paper.

E.1 The integral cover support lemma

In the context of Tetris-Reloaded, a box $\mathbf{a}$ is called an *input gap box* if it was loaded from $\mathcal{B}$ into $\mathcal{A}$ (i.e. if $\mathbf{a} \in \mathcal{B}'$) at some point in time during the execution of the algorithm. Note that we load into the knowledge base $\mathcal{A}$ either gap boxes from $\mathcal{B}$, or an output (unit) box. We distinguish between two types of resolvents. We call a resolvent $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ an *output resolvent* if either $\mathbf{w}_1$ or $\mathbf{w}_2$ is an output box, or (recursively) if either $\mathbf{w}_1$ or $\mathbf{w}_2$ is an output resolvent. Other resolvents are called *gap box resolvents*.

Definition E.2 (Projection of a box onto some support)**.** Let $\mathbf{b} = \langle x_1, \ldots, x_n \rangle$ be any dyadic box, and V be some subset of attributes. Then, the projection of $\mathbf{b}$ onto V, denoted by $\pi_V(\mathbf{b})$, is the box $\overline{\mathbf{b}} = \langle y_1, \ldots, y_n \rangle$ where

$$y_i = \begin{cases} x_i & \text{if } i \in V \\ \lambda & \text{if } i \notin V. \end{cases}$$

Definition E.3 (Resolvent supported on an integral cover)**.** Let $\mathbf{w}$ be a gap box resolvent. Let S be a subset of $\mathsf{support}(\mathbf{w})$, and $\mathbf{a}$ be an input gap box. Then, $\mathbf{w}$ is said to be *supported by* $\mathbf{a}$ *on* S if $\pi_S(\mathbf{a}) \subseteq \pi_S(\mathbf{w})$. (Geometrically, the shadow of $\mathbf{w}$ on the coordinate subspace of the variables in S contains the shadow of $\mathbf{a}$ on the same subspace.) An *integral cover* of $\mathsf{support}(\mathbf{w})$ is a collection of subsets of $\mathsf{support}(\mathbf{w})$, say $S_1, \ldots, S_c$, such that

$$S_1 \cup \cdots \cup S_c = \mathsf{support}(\mathbf{w}).$$

The resolvent $\mathbf{w}$ is said to be *supported on an integral cover* $S_1, \ldots, S_c$ if for each $i \in [c]$, $\mathbf{w}$ is supported by some input gap box on S_i.

Note that the collection $\{S_1, \ldots, S_c\}$ viewed as a hypergraph forms an integral (edge) cover of the ground set $\mathsf{support}(\mathbf{w})$.

Lemma E.4 (The integral cover support lemma)**.** *Suppose there is a positive integer $c \in [n]$ such that every gap box resolvent $\mathbf{w}$ is supported by an integral cover of size at most c. Then,* Tetris-Reloaded *runs in time $\tilde{O}(|C_\square|^c + Z)$.*

Proof. Noting Theorem C.5, we only need to show that the number of gap box resolvents is at most $\tilde{O}(|C_\square|^c)$. The total number of boxes that Tetris-Reloaded loads from $\mathcal{B}$ into $\mathcal{A}$ is $\sum_{\mathcal{B}'} |\mathcal{B}'|$. For a given subset K of the n attributes of $\mathcal{B}$, there are only $\tilde{O}(1)$ possible integral covers of size c. For each such integral cover $\{S_1, \ldots, S_c\}$, there are at most $\tilde{O}((\sum_{\mathcal{B}'} |\mathcal{B}'|)^c)$ witnesses $\mathbf{w}$ supported on this integral cover. (This is because given an integral cover of size c along with c input gap boxes , there are maximally $\tilde{O}(1)$ witnesses supported on this integral cover by those c input gap boxes, thanks to Proposition B.12.) By Lemma E.1, $(\sum_{\mathcal{B}'} |\mathcal{B}'|)^c = \tilde{O}(|C_\square|^c)$. The number of possible choices of K is also $\tilde{O}(1)$, which completes the proof. □

E.2 The $\tilde{O}(|C_\square|^{w+1} + Z)$ runtime for queries with treewidth w

In this section, we will prove Theorem 4.7. We begin with some background.

Definition E.5 ($\mathsf{support}(A_k)$ and the induced width of a GAO)**.** Let $\mathcal{H}_n = (\mathcal{V}, \mathcal{E})$ be a hypergraph whose vertex set is $\mathcal{V} = \{A_1, \ldots, A_n\}$. Let $\sigma = (A_1, \ldots, A_n)$ be a GAO for the vertices (attributes) in $\mathcal{V}$. The support of A_n (denoted by $\mathsf{support}(A_n)$) is the union of all hyperedges in $\mathcal{H}_n$ that contain A_n. Construct $\mathcal{H}_{n-1}$ from $\mathcal{H}_n$ by inserting $\mathsf{support}(A_n)$ as a new hyperedge, and then removing A_n from the vertex set and from all the hyperedges of $\mathcal{H}_n$. Then, define $\mathsf{support}(A_{n-1})$ as the union of all hyperedges containing A_{n-1} in $\mathcal{H}_{n-1}$. We keep constructing hypergraphs $\mathcal{H}_k$ and defining the supports of A_k, for $k = n-2, \ldots, 1$ in the same way:

- $\mathsf{support}(A_k)$ is the union of all hyperedges of $\mathcal{H}_k = (\mathcal{V}_k, \mathcal{E}_k)$ that contain A_k.
- $\mathcal{H}_{k-1}$ is constructed from $\mathcal{H}_k$ by adding a new hyperedge $\mathsf{support}(A_k)$ to $\mathcal{H}_k$ and removing the vertex A_k from $\mathcal{H}_k$.

The quantity

$$w = \max_{k \in [n]} |\mathsf{support}(A_k)| - 1 \tag{10}$$

is called the *induced width* of $\sigma = (A_1, \ldots, A_n)$ (with respect to the hypergraph $\mathcal{H}$). Furthermore, if $\mathcal{H}$ has treewidth w, then there exists a GAO with induced width w. This follows from the well-known fact that the smallest induced treewidth (over all elimination orders) of a hypergraph is the same as the treewidth of the hypergraph (see, e.g., [5,23]). Such a GAO with optimal induced width can be computed in time exponential in the size of $\mathcal{H}$. [20]

Notice that if $\mathcal{H}$ is α-acyclic and σ is a reversed GYO elimination order, then the above definition reduces back to Definition D.5.

Lemma E.6. *Let $\mathcal{B}$ be a set of boxes over the attributes $\{A_1, \ldots, A_n\}$ and suppose that we run* Tetris-Reloaded *on $\mathcal{B}$ with the SAO $\sigma = (A_1, \ldots, A_n)$. Let $\mathbf{w}$ be an input gap box or a gap box resolvent, and suppose that the last non-λ component of $\mathbf{w}$ is on attribute A_k (i.e. $k = \max\{i \mid A_i \in \mathsf{support}(\mathbf{w})\}$). Then*

$$\mathsf{support}(\mathbf{w}) \subseteq \mathsf{support}(A_k).$$

Proof. We prove this lemma by induction. If $\mathbf{w}$ is an input gap box, then clearly $\mathsf{support}(\mathbf{w}) \in \mathcal{E}_k$. For the inductive step, suppose $\mathbf{w}$ is the resolution of $\mathbf{w}_1$ and $\mathbf{w}_2$ on attribute A_k, where the induction hypothesis is $\mathsf{support}(\mathbf{w}_1) \subseteq \mathsf{support}(A_k)$ and $\mathsf{support}(\mathbf{w}_2) \subseteq \mathsf{support}(A_k)$.

If the last non-λ component of $\mathbf{w}$ is on A_k, then by Lemma C.1

$$\mathsf{support}(\mathbf{w}) \subseteq \mathsf{support}(\mathbf{w}_1) \cup \mathsf{support}(\mathbf{w}_2) \subseteq \mathsf{support}(A_k).$$

If the resolution turns the kth component of $\mathbf{w}$ into a λ, then $\mathsf{support}(\mathbf{w}) \subseteq \mathsf{support}(A_k) - \{A_k\}$. Suppose the last non-$\lambda$ component of $\mathbf{w}$ is on $A_{k'}$, then $A_{k'} \in \mathsf{support}(A_k) - \{A_k\}$, and hence $\mathsf{support}(\mathbf{w})$ is a hyperedge of the graph $\mathcal{H}_{k'}$. This means $\mathsf{support}(\mathbf{w}) \subseteq \mathsf{support}(A_{k'})$ as desired. □

For the sake of completeness, we re-state Theorem 4.7 and then prove it.

Theorem E.7 (Theorem 4.7 re-stated)**.** *For any set of boxes $\mathcal{B}$ with $\mathsf{tw}(\mathcal{H}(\mathcal{B})) = w$, by setting SAO to be the attribute ordering with elimination width w,* Tetris-Reloaded *solves* BCP *on input $\mathcal{B}$ in time $\tilde{O}(|\mathcal{C}_\square|^{w+1} + Z)$. Here, $\mathcal{C}_\square$ is any optimal box certificate for the instance, and Z is the output size.*

Proof. We apply Lemma E.4. We will show that for every witness $\mathbf{w}$ which is a either an input gap box or a gap box resolvent, its support $\mathsf{support}(\mathbf{w})$ is the union of $w+1$ singleton sets $V_1, \ldots, V_{w+1}$ such that for each set V_i, $\pi_{V_i}(\mathbf{a}) \subseteq \pi_{V_i}(\mathbf{w})$ for some input gap box $\mathbf{a}$.

From Lemma E.6 and equation (10), it follows that $\mathsf{support}(\mathbf{w}) \leqslant w+1$. Therefore, it is sufficient to prove the following claim: every non-λ component of $\mathbf{w}$ is a prefix of some component of an input gap box $\mathbf{a}$.

The claim is proved easily by induction. For the base case, if $\mathbf{w}$ is an input gap box, then clearly the claim holds. For the inductive step, suppose $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ where the claim holds for $\mathbf{w}_1$ and $\mathbf{w}_2$, and the resolution is on attribute A_k. The claim holds for $\mathbf{w}$ because every non-λ component of $\mathbf{w}$ is either the same as that of $\mathbf{w}_1$ or of $\mathbf{w}_2$, except for the component corresponding to A_k which is a prefix of the component from $\mathbf{w}_1$ (and $\mathbf{w}_2$). □

[20] If $\mathcal{H}$ is a hypergraph representing a join query, then the size of $\mathcal{H}$ is data-independent. We say that $\mathcal{H}$ has size $\tilde{O}(1)$ in data complexity.

E.3 The $\tilde{O}(|C_\square| + Z)$ runtime for queries with treewidth 1

Given a hypergraph $\mathcal{H}$ (or a query Q), there exists a GAO with induced width 1 if and only if $\mathcal{H}$ is a forest (all relations in Q have arity at most 2, and Q is α-acyclic [21]). When the induced width is 1, by definition $\mathsf{support}(A_k) \leqslant 2$ for all $k \in [n]$.

The essence of the proof of Theorem 4.5 is to make use of the fact that when resolving $\mathbf{w}_1$ and $\mathbf{w}_2$ with support of size at most 2, we end up with a resolvent $\mathbf{w}$ which contains either $\mathbf{w}_1$ or $\mathbf{w}_2$ or both. (See Figure 6 for an illustration.)

Lemma E.8 (2D-resolution expands)**.** *Suppose* Tetris *performs* $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ *such that* $|\mathsf{support}(\mathbf{w}_1) \cup \mathsf{support}(\mathbf{w}_2)| \leqslant 2$*, then the resulting box* $\mathbf{w}$ *is a prefix box* [22] *of either* $\mathbf{w}_1$ *or* $\mathbf{w}_2$ *or both.*

Proof. Given two strings x_1, x_2 and two boxes $\mathbf{w_1}, \mathbf{w_2}$ WLOG of the form:

$$\begin{aligned} \mathbf{w}_1 &= \langle \lambda\ ,\ \ x_1\ \ ,\ \lambda\ ,\ \lambda\ ,\ x_2 0\ ,\ \lambda \rangle \\ \mathbf{w}_2 &= \langle \lambda\ ,\ \text{prefix}\,(x_1)\ ,\ \lambda\ ,\ \lambda\ ,\ x_2 1\ ,\ \lambda \rangle \end{aligned}$$

The resolvent of $\mathbf{w}_1$ and $\mathbf{w}_2$ is

$$\mathbf{w} = \langle \lambda\ ,\ x_1\ ,\ \lambda\ ,\ \lambda\ ,\ x_2\ ,\ \lambda \rangle.$$

By definition C.2, $\mathbf{w}$ is a prefix box of $\mathbf{w}_1$. □

Theorem E.9 (Theorem 4.5 re-stated)**.** *For any set of boxes* $\mathcal{B}$ *with* $\mathsf{tw}(\mathcal{H}(\mathcal{B})) = 1$*, by setting SAO to be the attribute ordering with elimination width* 1*,* Tetris-Reloaded *solves* BCP *on input* $\mathcal{B}$ *in time* $\tilde{O}(|C_\square| + Z)$*. Here,* $C_\square$ *is any optimal box certificate for the instance, and* Z *is the output size.*

Proof. In light of Lemma E.4, we show that every gap box resolvent $\mathbf{w}$ is supported by *one* input gap box on $\mathsf{support}(\mathbf{w})$. In particular, we use induction to show that $\mathbf{a} \subseteq \mathbf{w}$ for some input gap box $\mathbf{a}$. If $\mathbf{w}$ was an input gap box, then the above obviously holds. We can use that as a base case.

Tetris selects a SAO with elimination width 1. From Lemma E.6 and Equation 10, any boxes $\mathbf{w}_1$ and $\mathbf{w}_2$ that are resolved (on any dimension) satisfy $|\mathsf{support}(\mathbf{w}_1) \cup \mathsf{support}(\mathbf{w}_2)| \leqslant 2$. From Lemma E.8, every resolution $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ results in a box $\mathbf{w}$ containing either $\mathbf{w}_1$ or $\mathbf{w}_2$. Hence, by induction every gap box resolvent contains an input gap box. □

E.4 The $\tilde{O}(|C_\square|^{n-1} + Z)$ runtime for any query over n attributes

In this section, we will prove the following result. (An n-clique is an example of a hypergraph whose treewidth $w = n - 1$.)

Theorem E.10 (Tetris-Reloaded runs in time $\tilde{O}(|C_\square|^{n-1}+Z)$)**.** *For any set of boxes* $\mathcal{B}$ *in* n *dimensions,* Tetris-Reloaded *solves* BCP *on input* $\mathcal{B}$ *in time* $\tilde{O}(|C_\square|^{n-1} + Z)$*. Here,* $C_\square$ *is any optimal box certificate for the instance, and* Z *is the output size.*

Along with Proposition 3.6, the above result implies the following:

Corollary E.11. Tetris-Reloaded *evaluates any join query* Q *over* n *attributes in time* $\tilde{O}(|C_\square|^{n-1} + Z)$*, where* $C_\square$ *is an optimal box certificate for the join instance, and* Z *is the output size.*

Proof of Theorem E.10. We inductively prove the following claim: Every witness $\mathbf{w}$ which is involved in a gap box resolution on A_n is supported by one input gap box on $\{A_i, A_n\}$ for each $i \in [n-1]$. In the base case, $\mathbf{w}$ is an input gap box in which case the claim is obviously true. For the inductive step, $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$. For each $i \in [n-1]$, either $\pi_{\{A_i,A_n\}}(\mathbf{w}_1) \subseteq \pi_{\{A_i,A_n\}}(\mathbf{w})$ or $\pi_{\{A_i,A_n\}}(\mathbf{w}_2) \subseteq \pi_{\{A_i,A_n\}}(\mathbf{w})$. Assuming the claim holds for $\mathbf{w}_1$ and $\mathbf{w}_2$, it holds for $\mathbf{w}$.

For each $i \in [n-1]$, every witness $\mathbf{w}$ which is involved in a gap box resolution on A_i is supported on its support set by $\leqslant i$ input gap boxes. Lemma E.4 completes the proof. □

[21] When all relations have arity $\leqslant 2$, α-acyclicity and β-acyclicity coincide.

[22] See Definition C.2 for the notion of "prefix box". Note that in this situation $\mathbf{w}$ being a prefix box of $\mathbf{w}_i$ is the same as $\mathbf{w}$ containing $\mathbf{w}_i$.

F Omitted details from Section 4.5

F.1 The high-level ideas

While Theorem E.10 shows that Tetris-Reloaded runs in time $\tilde{O}(|C_\square|^{n-1}+Z)$ for n-attributes/dimensions, Theorem 5.4 shows a lower bound of $\Omega(|C_\square|^{n-1}+Z)$ for all algorithms that use only Ordered Geometric Resolution.

In this section, we highlight the key ideas that lead to a new enhancement of Tetris that enables it to bypass the above lower bound and solve join queries (and more generally BCP instances) in time $\tilde{O}(|C_\square|^{n/2}+Z)$. We take two specific values of n as examples: $n=3$ and $n=4$, attempting to introduce one new idea at a time. Note that this result is independent of the input query structure.

In these examples, we assume that the algorithm is given as its input the box certificate of the instance of BCP, which we will denote by C. We will also refer to this version of BCP as the *offline* case of the problem. Moreover, we will only consider the *Boolean* version of BCP where the objective is to determine whether the output is empty. Boolean BCP is equivalent to determining whether the set C of input gap boxes covers the entire output space. (See Definition 3.5.)

Later on in Section F.5, we will be using Theorem C.5 to convert the offline algorithm solving *Boolean* BCP in time $\tilde{O}(|\mathcal{C}|^{n/2})$ into an offline algorithm solving BCP in time $\tilde{O}(|\mathcal{C}|^{n/2}+Z)$. In Section F.6, we will show how to convert the *offline* algorithm into an *online* one. Theorem 5.5 shows a lower bound of $\Omega(|\mathcal{C}|^{n/2}+Z)$ for all Geometric Resolution algorithms.

F.2 Idea 1: load-balancing

We start with a motivating example.

Example F.1. Consider an input query Q over $n=3$ attributes X, Y, and W, and the set $C := C_1 \cup C_2 \cup C_3$ of input gap boxes, where

$$\begin{aligned}
C_1 &= \left\langle 0x,\lambda,0\right\rangle \;\middle|\; x\in\{0,1\}^{d-2}\Big\} \;\cup\; \Big\{\langle 0,y,1\rangle \;\middle|\; y\in\{0,1\}^{d-2}\right\}\\
C_2 &= \left\langle 10x,0,\lambda\right\rangle \;\middle|\; x\in\{0,1\}^{d-2}\Big\} \;\cup\; \Big\{\langle 10,1,z\rangle \;\middle|\; z\in\{0,1\}^{d-2}\right\}\\
C_3 &= \left\langle 110,y,\lambda\right\rangle \;\middle|\; y\in\{0,1\}^{d-2}\Big\} \;\cup\; \Big\{\langle 111,\lambda,z\rangle \;\middle|\; z\in\{0,1\}^{d-2}\right\}.
\end{aligned}$$

Note that C_1, C_2, C_3, and C cover $\langle 0,\lambda,\lambda\rangle$, $\langle 10,\lambda,\lambda\rangle$, $\langle 11,\lambda,\lambda\rangle$, and $\langle\lambda,\lambda,\lambda\rangle$ respectively. The output is empty and only $O(|C|)$ out-of-order geometric resolutions are sufficient to prove it.[23] However, $\Omega(|C|^2)$ ordered geometric resolutions are necessary for the proof of $\langle\lambda,\lambda,\lambda\rangle$, no matter what the SAO is. This is because if W is the last attribute in the SAO, then we will need $\Omega(|C|^2)$ resolutions just to infer $\langle 0,\lambda,\lambda\rangle$ from C_1. If Y (or X) is the last attribute, then C_2 (or C_3) is going to create the same problem.

The main reason Tetris-Preloaded (and any Ordered Geometric Resolution algorithm) is slow in the kind of inputs shown above is because it got "stuck" in performing two many resolutions on one particular attribute, creating many witnesses along a particular dimension. At a very high level, our first idea is to cut the output space along a particular dimension into a collection of "layers," each of which has relatively "few" input gap boxes. Then, our algorithm explores whether each of these layers is completely covered by the input gap boxes. Since those layers form a partition of a given dimension, the output space is covered iff each layer is covered. The layers are chosen so that the number of layers is "small" and the number of boxes that can contribute to covering each layer is also "small." From there, verifying that a layer is covered takes little time.

To be more concrete, we first define the notion of "dimension partition."

Definition F.2 (Dimension partition)**.** Recall that the domain of each attribute is assumed to be $\mathbf{D}=\{0,1\}^d$, i.e. the set of all binary strings of length d. Each binary string x of length $\leqslant d$ is a dyadic interval. A *partition* P of $\mathbf{D}=\{0,1\}^d$ is

[23] Indeed note that by resolving on the X-attribute of the first subset of C_1, we can get $\langle 0,\lambda,0\rangle$. Similarly by resolving on the Y-attribute of the second subset of C_1, we get the box $\langle 0,\lambda,1\rangle$. With one more resolution we get the box $\langle 0,\lambda,\lambda\rangle$ from C_1. Similarly by doing resolutions on X, W, then Y-attributes of C_2, one can get the box $\langle 10,\lambda,\lambda\rangle$. Finally, by resolving on Y, W, then X attributes of C_3, we can get the box $\langle 11,\lambda,\lambda\rangle$. Then with two more resolutions on the three boxes from C_1, C_2, C_3, we get $\langle\lambda,\lambda,\lambda\rangle$, as desired.

a collection of disjoint dyadic intervals whose union is exactly $\mathbf{D}$. In particular, P is a partition of $\mathbf{D}$ iff it satisfies the following two properties:

- Strings in P are prefix-free, i.e. no string is a prefix of another.
- For every string s of length d, there is some string $x \in P$ so that x is a prefix of s.

Given an attribute X of the input query, an *X-partition* is a partition of the domain $\mathbf{D}(X)$. We will typically use P_X to denote a partition along dimension X.

Example (F.1 Continued). $P_X = \{0, 10, 110, 111\}$ is one possible X-partition. Another X-partition could be $P_X = \{0, 100, 101, 110, 111\}$.

Geometrically, a partition along dimension X divides the output space into $|P_X|$ "layers", one for each interval x in P_X. In particular, the layer defined by a fixed interval $x \in P_X$, called the x-layer, is the dyadic box $\langle x, \lambda, \lambda \rangle$. To verify that the output space is completely covered by the input gap boxes, it is sufficient to verify that every x-layer is covered, for each $x \in P_X$. An input gap box whose X-component is disjoint from x will not affect whether the x-layer is covered. Hence, to verify whether the x-layer is covered, we can ignore all gap boxes that do not intersect x. If the remaining set of gap boxes is small, then this verification is fast. At the same time, we do not want too many layers because that certainly increases the total amount of verification work. This balancing act leads to our first idea: we find a dimension partition that is somehow "balanced."

For notational convenience, given a set of boxes C and a dyadic interval x on the domain $\mathbf{D}(X)$, define two sets:

$$C_{\subset x}(X) = \{\mathbf{b} \in C \mid \pi_X(\mathbf{b}) \subsetneq x\} \tag{11}$$

$$C_{\supseteq x}(X) = \{\mathbf{b} \in C \mid \pi_X(\mathbf{b}) \supseteq x\}. \tag{12}$$

In other words, the first set $C_{\subset x}(X)$ consists of all dyadic boxes in C that are strictly contained in the x-layer. The second set $C_{\supseteq x}(X)$ is the set of boxes in C each of whose X-component completely covers the interval x. Note that, for every box $\mathbf{b} \in C - (C_{\subset x}(X) \cup C_{\supseteq x}(X))$, the dyadic interval $\pi_X(\mathbf{b})$ is completely disjoint from the interval x. It should also be emphasized that $\pi_X(\mathbf{b}) \subsetneq x$ means x (as a string) is a strict prefix of the binary string $\pi_X(\mathbf{b})$. (Alternatively, x as a dyadic segment strictly contains the dyadic segment $\pi_X(\mathbf{b})$.)

Definition F.3 (Balanced dimension partition)**.** Let C be the set of input gap boxes, and X be any attribute. A *balanced X-partition* is an X-partition P_X such that

$$|P_X| = \tilde{O}(\sqrt{|C|})$$
$$|C_{\subset x}(X)| \leqslant \sqrt{|C|}, \text{ for every } x \in P_X.$$

Geometrically, the first condition states that the number of layers is small. The second condition states that the number of input gap boxes completely contained in every x-layer is small. Since there are $|C|$ input gap boxes, the quantity $\sqrt{|C|}$ is exactly the mid-point, balancing the number of layers and the number of boxes contained within each layer.

Example (F.1 Continued). The following is a balanced X-partition: (Recall that $|C| = 6 \cdot 2^{d-2}$ and that Definition F.3 does not put any restrictions on $|C_{\supseteq x}(X)|$.)

$$P_X = \left\{0x' \mid x' \in \{0,1\}^{\lceil (d-2)/2 \rceil}\right\} \cup \left\{10x' \mid x' \in \{0,1\}^{\lceil (d-2)/2 \rceil}\right\} \cup \left\{110, 111\right\}. \tag{13}$$

A balanced partition can be constructed easily, as the following proposition shows.

Proposition F.4 (Efficient constructibility of balanced partitions)**.** *Given a set C of input gap boxes and an arbitrary attribute X, a balanced X-partition can be computed in time $\tilde{O}(|C|)$.*

Proof. We start with the trivial partition $P = \{\lambda\}$ that has only one layer, and keep revising it until it becomes balanced. An interval $x \in P$ is said to be *heavy* if

$$\left|C_{\subset x}(X)\right| > \sqrt{|C|}. \tag{14}$$

While there is still a heavy interval x in P, replace x by the two sub-intervals $x0, x1$. These sub-intervals are called *children* of x, and x is a *parent* of both $x0$ and $x1$. This process certainly terminates when no interval in P is heavy anymore. (Note that by definition a unit interval x is not heavy because $\left|C_{\subset x}(X)\right|$=0.) It remains to show that $|P| = \tilde{O}(\sqrt{|C|})$ in the end. To see this, consider the set H of all heavy intervals. If an interval is heavy then all its prefixes are heavy. Let H' be the maximal set of prefix-free intervals in H. Then the intervals in H' are disjoint. Because they are all heavy, $|H'| \leqslant \sqrt{|C|}$. Furthermore, because the intervals in $H - H'$ are prefixes of intervals in H', we conclude that $|H| = \tilde{O}(\sqrt{|C|})$. Finally, since the intervals in P are all children of some intervals in H, we conclude that $|P| \leqslant 2|H| = \tilde{O}(\sqrt{|C|})$ as well. □

Next, we explain how a balanced X-partition can be used to solve the offline version of the problem in time $\tilde{O}(|C|^{3/2})$ when $n = 3$. Consider an input query Q over $n = 3$ attributes X, Y, and W. Let P_X be a balanced X-partition. For every interval $x \in P_X$, we would like to verify that the x-layer, i.e. the box $\langle x, \lambda, \lambda\rangle$ is covered by all gap boxes in C. The boxes in $C_{\subset x}(X) \cup C_{\supseteq x}(X)$ are the only boxes in C that intersect the x-layer. Hence, it is sufficient to verify that boxes in this union cover the x-layer.

At this point, we introduce another simple idea: we reduce the x-layer coverage problem above to a slightly different 3-dimensional coverage problem and run Tetris-Preloaded on it. Define

$$C[x] = \left\{\langle \lambda, y, w\rangle \mid \langle x', y, w\rangle \in C_{\supseteq x}(X) \text{ for some } x'\right\} \tag{15}$$

$$\mathcal{F}[x] = \left\{\langle x', y, w\rangle \mid \langle xx', y, w\rangle \in C_{\subset x}(X)\right\}. \tag{16}$$

Then, verifying whether the boxes in $C_{\subset x}(X) \cup C_{\supseteq x}(X)$ cover the x-layer is exactly identical to verifying whether $C[x] \cup \mathcal{F}[x]$ covers the entire space $\langle \lambda, \lambda, \lambda\rangle$. In other words, because the X-components of the boxes in $C_{\supseteq x}(X)$ already contain the entire x-interval, we might as well truncate all but the x-interval part of those boxes. And, we convert the X-component of boxes in $C_{\subset x}(X)$ to "relative" values in the coordinate system restricted to the x-layer.

Now, we run Tetris-Preloaded on the $C[x] \cup \mathcal{F}[x]$ input with the SAO being (Y, W, X). While this SAO might seem a bit unnatural, we will show in the next sub-section why this SAO embodies a new idea. We bound the runtime of Tetris-Preloaded using Lemma E.4.

- If a witness $\mathbf{w}$ is a result of a gap box resolution on X, then its support is (integrally) covered by two boxes from $\mathcal{F}[x]$ (as all boxes from $C[x]$ has X-component already equal λ). Thus, the number of such witnesses is at most $\tilde{O}(|\mathcal{F}[x]|^2) = \tilde{O}(\left|C_{\subset x}(X)\right|^2) = \tilde{O}(|C|)$.
- The number of witnesses $\mathbf{w}$ resulting from gap box resolutions on Y or W is also at most $|C|$. For example, consider a witness $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ which is resolved on Y or W. By Lemma E.8, $\mathbf{w}$ is a prefix box of either $\mathbf{w}_1$ or $\mathbf{w}_2$. The witnesses $\mathbf{w}_i$ ($i \in \{1, 2\}$) fall into three classes: (1) $\mathbf{w}_i$ is the result of a gap box resolution on Y or W, (2) $\mathbf{w}_i \in C[x]$, (3) $\mathbf{w}_i$ is the result of gap box resolution on X (the resolution turns $\mathbf{w}_i$'s X-component into a λ). The number of witnesses of types (2) and (3) is at most $\tilde{O}(|C|)$. Hence, by induction the number of gap box resolutions on Y or W is also $\tilde{O}(|C|)$.

(We will make the above argument more formal in the next section.) Hence, the overall x-layer verification process takes $\tilde{O}(|C|)$ time; and, since there are $\tilde{O}(\sqrt{|C|})$ intervals in P_X, the entire algorithm takes $\tilde{O}(|C|^{3/2})$-time as desired.

Example. F.1 Continued We apply (16) and (15) for every x in the X-balanced partition P_X given earlier in (13). For every $x' \in \{0,1\}^{\lceil (d-2)/2\rceil}$, we get (assuming $d \geqslant 4$):

$$\begin{aligned}
\mathcal{F}[0x'] &= \left\{\langle x'', \lambda, 0\rangle \mid x'' \in \{0,1\}^{\lfloor (d-2)/2\rfloor}\right\} \\
C[0x'] &= \left\{\langle \lambda, y, 1\rangle \mid y \in \{0,1\}^{d-2}\right\} \\
\mathcal{F}[10x'] &= \left\{\langle x'', 0, \lambda\rangle \mid x'' \in \{0,1\}^{\lfloor (d-2)/2\rfloor}\right\} \\
C[10x'] &= \left\{\langle \lambda, 1, z\rangle \mid z \in \{0,1\}^{d-2}\right\}.
\end{aligned}$$

Moreover

$$\begin{aligned} C[110] &= \left\{\langle \lambda, y, \lambda\rangle \quad | \quad y \in \{0,1\}^{d-2}\right\} \\ C[111] &= \left\{\langle \lambda, \lambda, z\rangle \quad | \quad z \in \{0,1\}^{d-2}\right\} \\ \mathcal{F}[110] &= \mathcal{F}[111] = \varnothing. \end{aligned}$$

Now for every $x' \in \{0,1\}^{\lceil (d-2)/2 \rceil}$, we verify that the $0x'$-layer is covered by verifying that $C[0x'] \cup \mathcal{F}[0x']$ covers $\langle \lambda, \lambda, \lambda\rangle$. While running Tetris-Preloaded on $C[0x'] \cup \mathcal{F}[0x']$ using the SAO of (Y, W, X), we need $\tilde{O}(\sqrt{|C|})$ resolutions on X to infer $\langle \lambda, \lambda, 0\rangle$ from $\mathcal{F}[0x']$. After that, we need $\tilde{O}(|C|)$ resolutions on Y and W to infer $\langle \lambda, \lambda, \lambda\rangle$ from $C[0x'] \cup \{\langle \lambda, \lambda, 0\rangle\}$. (Recall that we were looking at the $0x'$-layer.)

In a similar way, we can handle the remaining $10x', 110, 111$-layers.

F.3 Idea 2: taking the high road (to higher dimension)

The algorithm described above can be expressed more cleanly as follows. Strategies based on ordered geometric resolutions do not meet the runtime target (of $\tilde{O}(|C|^{n/2})$) for the original problem. However, it will meet the runtime target if we map each input gap box to a higher dimensional gap box, *then* run Tetris-Preloaded on these new gap boxes. This idea is similar in spirit to the kernel method in machine learning: data in the original dimensions are not linearly separable, but they become linearly separable in higher dimensions after a kernel map.

More concretely, we will re-cast the above algorithm by explicitly constructing the map that transforms each input gap box to a gap box in a higher-dimensional space. This set of new gap boxes are constructed so that they are "load-balanced", making Tetris work efficiently on them.

Let s be any dyadic interval, and let $\text{prefixes}(s)$ denote the set of all binary strings that are prefixes of s, including s itself. For any set S of dyadic intervals. Define

$$\text{prefixes}(S) = \bigcup_{s \in S} \text{prefixes}(s).$$

Let P_X denote a fixed balanced X-partition of C. Let X' and X'' be two new attribute names. The following map, called the *X-load balancing map*,

$$\mathsf{Balance}_X : \mathbf{D}(X) \times \mathbf{D}(Y) \times \mathbf{D}(W) \to \mathbf{D}(X') \times \mathbf{D}(Y) \times \mathbf{D}(W) \times \mathbf{D}(X''),$$

is defined as follows. Let $\mathbf{b} = \langle x, y, w\rangle$ be any box in $\mathbf{D}(X) \times \mathbf{D}(Y) \times \mathbf{D}(W)$. Then,

$$\mathsf{Balance}_X(\mathbf{b}) = \mathsf{Balance}_X(\langle x, y, w\rangle) := \begin{cases} \langle x', y, w, x''\rangle & \text{if } x = x'x'' \text{ where } x' \in P_X \text{ and } x'' \neq \lambda \\ \langle x, y, w, \lambda\rangle & \text{otherwise. (i.e., } x \in \text{prefixes}(P_X)) \end{cases}.$$

Let $\mathsf{Balance}_X(C) = \left\{\mathsf{Balance}_X(\mathbf{b}) \mid \mathbf{b} \in C\right\}$. Then the algorithm described in the previous section is simply to run Tetris-Preloaded on input $\mathsf{Balance}_X(C)$. This input has gap boxes in 4 dimensions (X', Y, W, X''), which is also the SAO the algorithm adopts.

For each $x \in P_X$, define the following two sets:

$$C[x] = \left\{\langle x', y, w, \lambda\rangle \in \mathsf{Balance}_X(C) \mid x' \in \text{prefixes}(x)\right\} \tag{17}$$

$$\mathcal{F}[x] = \left\{\langle x, y, w, x''\rangle \in \mathsf{Balance}_X(C) \mid x'' \neq \lambda\right\}. \tag{18}$$

Note that

$$\mathsf{Balance}_X(C) = \bigcup_{x \in P_X} (C[x] \cup \mathcal{F}[x]),$$

and that the sets $\mathcal{F}[x]$ are disjoint for $x \in P_X$. Because P_X is a balanced X-partition, we know $|\mathcal{F}[x]| \leqslant \sqrt{|C|}$ for all $x \in P_X$ and, obviously, $|C[x]| \leqslant |C|$.

Definition F.5 (A-witness)**.** Let A be an attribute from the set $\{X', Y, W, X''\}$. A witness $\mathbf{w}$ is called an *A-witness* if it participates in a gap box resolution on attribute A during the execution of the algorithm. (If $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$, then all three boxes $\mathbf{w}, \mathbf{w}_1, \mathbf{w}_2$ participate in the resolution.)

We analyze the above algorithm by counting the number of A-witnesses for each $A \in \{X', Y, W, X''\}$.

- First, we bound the number of X''-witnesses. When Tetris resolves two boxes $\mathbf{w}_i = \langle x'_i, y_i, w_i, x''_i \rangle, i \in \{1, 2\}$ on X'', the strings x'_1 and x'_2 must be a prefix of one another. Furthermore, $x''_i \neq \lambda$ for $i \in \{1, 2\}$. Hence, $x'_1, x'_2 \in P_X$. But because strings in P_X are prefix-free, $x'_1 = x'_2$. Consequently, the gap box resolutions on X'' can be grouped into disjoint groups, one for each $x \in P_X$. For each such $x \in P_X$, the number of X''-witnesses $\mathbf{w} = \langle x, y, w, x'' \rangle$ is at most $\tilde{O}(|\mathcal{F}[x]|^2)$ because $\mathbf{w}$ must be supported on $\{Y, X''\}$ and $\{W, X''\}$ by boxes from $\mathcal{F}[x]$. Consequently, the total number of X''-witnesses is at most $\sum_{x \in P_X} \tilde{O}(|\mathcal{F}[x]|^2) = \tilde{O}(|C|^{3/2})$.

- Similarly, we count the number of W-witnesses by fixing an $x \in P_X$ and counting the number of W-witnesses of the form $\mathbf{w} = \langle x', y, w, \lambda \rangle$ for some string $x' \in \text{prefixes}(x)$. Then, by summing this count over all $x \in P_X$, we obtain an upperbound on the number of W-witnesses. We sketch the counting argument below, leaving the completely rigorous description to the next sections when we describe the general algorithm and its formal analysis.

  Suppose $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$ is a W-witness. Note that since Tetris only performs ordered resolutions, $\pi_{X''}(\mathbf{w}_1) = \pi_{X''}(\mathbf{w}_2) = \lambda$, and $\pi_W(\mathbf{w}_1) \neq \lambda, \pi_W(\mathbf{w}_2) \neq \lambda$. By Lemma E.8, $\pi_{Y,W}(\mathbf{w})$ is a prefix box of either $\pi_{Y,W}(\mathbf{w}_1)$ or $\pi_{Y,W}(\mathbf{w}_2)$. Each witness $\mathbf{w}_i$ ($i \in \{1, 2\}$) belongs to one of three classes: (1) $\mathbf{w}_i$ is the result of a gap box resolution on W, (2) $\mathbf{w}_i \in C[x]$, (3) $\mathbf{w}_i$ is the result of gap box resolution on X'' (this resolution makes $\pi_{X''}(\mathbf{w}_i) = \lambda$). The number of witnesses of types (2) and (3) is at most $\tilde{O}(|C|)$. (For (3), this is because $\mathbf{w}_i$ is supported on $\{Y, X''\}$ and $\{W, X''\}$ by two boxes from $\mathcal{F}[x]$.) Hence, for each $x \in P_X$, by induction the number of W-witnesses of the form $\langle x', y, w, \lambda \rangle$ where $x' \in \text{prefixes}(x)$ is $\tilde{O}(|C|)$. Overall, the total number of W-witnesses is $\tilde{O}(|C|^{3/2})$ as desired.

- The cases for Y and X' are much simpler and thus omitted.

Example. F.1 Continued We have $\mathsf{Balance}_X(C) = \mathsf{Balance}_X(C_1) \cup \mathsf{Balance}_X(C_2) \cup \mathsf{Balance}_X(C_3)$, where

$$\begin{aligned}
\mathsf{Balance}_X(C_1) &= \left\{\langle 0x', \lambda, 0, x'' \rangle \;\middle|\; x' \in \{0,1\}^{\lceil (d-2)/2 \rceil}, x'' \in \{0,1\}^{\lfloor (d-2)/2 \rfloor}\right\} \cup \\
&\quad \left\{\langle 0, y, 1, \lambda \rangle \;\middle|\; y \in \{0,1\}^{d-2}\right\} \\
\mathsf{Balance}_X(C_2) &= \left\{\langle 10x', 0, \lambda, x'' \rangle \;\middle|\; x' \in \{0,1\}^{\lceil (d-2)/2 \rceil}, x'' \in \{0,1\}^{\lfloor (d-2)/2 \rfloor}\right\} \cup \\
&\quad \left\{\langle 10, 1, z, \lambda \rangle \;\middle|\; z \in \{0,1\}^{d-2}\right\} \\
\mathsf{Balance}_X(C_3) &= \left\{\langle 110, y, \lambda, \lambda \rangle \;\middle|\; y \in \{0,1\}^{d-2}\right\} \cup \\
&\quad \left\{\langle 111, \lambda, z, \lambda \rangle \;\middle|\; z \in \{0,1\}^{d-2}\right\}.
\end{aligned}$$

Applying (18) and (17) on P_X from (13), for any $x' \in \{0,1\}^{\lceil (d-2)/2 \rceil}$, we get (assuming $d \geqslant 4$)

$$\begin{aligned}
\mathcal{F}[0x'] &= \left\{\langle 0x', \lambda, 0, x'' \rangle \;\middle|\; x'' \in \{0,1\}^{\lfloor (d-2)/2 \rfloor}\right\} \\
C[0x'] &= \left\{\langle 0, y, 1, \lambda \rangle \;\middle|\; y \in \{0,1\}^{d-2}\right\} \\
\mathcal{F}[10x'] &= \left\{\langle 10x', 0, \lambda, x'' \rangle \;\middle|\; x'' \in \{0,1\}^{\lfloor (d-2)/2 \rfloor}\right\} \\
C[10x'] &= \left\{\langle 10, 1, z, \lambda \rangle \;\middle|\; z \in \{0,1\}^{d-2}\right\}.
\end{aligned}$$

Moreover

$$\begin{aligned}
C[110] &= \left\{\langle 110, y, \lambda, \lambda \rangle \;\middle|\; y \in \{0,1\}^{d-2}\right\} \\
C[111] &= \left\{\langle 111, \lambda, z, \lambda \rangle \;\middle|\; z \in \{0,1\}^{d-2}\right\} \\
\mathcal{F}[110] &= \mathcal{F}[111] = \varnothing.
\end{aligned}$$

F.4 Idea 3: global accounting

We next consider a query over $n = 4$ attributes. We explain why a straightforward generalization of the balanced partition idea above does not quite work for $n \geqslant 4$. Then, we introduce a simple "global accounting" idea that makes it work.

Consider a query Q on attributes X, Y, W, T. Let C be the set of input gap boxes. As before, let P_X denote a balanced X-partition, and define the *X-load balancing map*

$$\mathsf{Balance}_X : \mathbf{D}(X) \times \mathbf{D}(Y) \times \mathbf{D}(W) \times \mathbf{D}(T) \to \mathbf{D}(X') \times \mathbf{D}(Y) \times \mathbf{D}(W) \times \mathbf{D}(T) \times \mathbf{D}(X''),$$

by setting

$$\mathsf{Balance}_X(\langle x, y, w, t\rangle) := \begin{cases} \langle x', y, w, t, x''\rangle & \text{if } x = x'x'' \text{ where } x' \in P_X \text{ and } x'' \neq \lambda \\ \langle x, y, w, t, \lambda\rangle & \text{otherwise. (i.e., } x \in \text{prefixes}\,(P_X)) \end{cases}.$$

Let $\mathsf{Balance}_X(C) = \Big\{\mathsf{Balance}_X(\mathbf{b}) \mid \mathbf{b} \in C\Big\}$. Similar to the previous section, also define

$$\begin{aligned} C[x] &= \Big\{\langle x', y, w, t, \lambda\rangle \in \mathsf{Balance}_X(C) \mid x' \in \text{prefixes}\,(x)\Big\} \\ \mathcal{F}[x] &= \Big\{\langle x, y, w, t, x''\rangle \in \mathsf{Balance}_X(C) \mid x'' \neq \lambda\Big\}. \end{aligned}$$

The obvious generalization is to run Tetris-Preloaded with $\mathsf{Balance}_X(C)$ as the input, and (X', Y, W, T, X'') as the SAO. We next attempt to count the number of A-witnesses for $A \in \{X', Y, W, T, X''\}$.

- Noting that $|\mathcal{F}[x]|| \leqslant \sqrt{|C|}$ for each $x \in P_X$, the number of X''-witnesses satisfying $\pi_{X''}(\mathbf{w}) = x$ is at most $\tilde{O}(\sqrt{|C|^3}) = \tilde{O}(|C|^{3/2})$; because each such witness is supported by (at most) three gap boxes from $\mathcal{F}[x]$, on $\{X', X'', Y\}$, $\{X', X'', W\}$, and $\{X', X'', T\}$. (We can show this easily by induction.) Along with the fact that there are $\tilde{O}(\sqrt{|C|})$ intervals $x \in P_X$ which are prefix-free, the total number of X''-witnesses is within the desired bound of $\tilde{O}(|C|^2)$.
- We next count the number of T-witnesses. This is where there is a problem. A T-witness $\mathbf{w}$ might be $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$, where $\mathbf{w}_1, \mathbf{w}_2 \in C[x]$, and the size of $C[x]$ might be up to $\tilde{\Omega}(|C|)$. The number of such witnesses $\mathbf{w}$ satisfying $\pi_{X'}(\mathbf{w}) \in \text{prefixes}\,(x)$ might be as many as $\tilde{\Omega}(|C|^2)$. Multiplying by $|P_X|$ and we get an $\tilde{\Omega}(|C|^{5/2})$ bound: too large! The problem here is the same as the $n = 3$ case when no balanced partition was introduced.

Hence, our next natural idea is to use a balanced Y-partition too, and break the x-layer coverage problem into multiple subproblems, one for each interval y in the balanced Y-partition of $C[x]$.

Before defining a more refined notion of balanced map, we need a few terminologies. Let P be a domain partition. Let s be an arbitrary dyadic interval. Define

$$s'(P) = \begin{cases} s & \text{if } s \in \text{prefixes}\,(P) \\ s' & \text{if } s = s's'' \text{ for some } s' \in P. \end{cases} \tag{19}$$

and

$$s''(P) = \begin{cases} \lambda & \text{if } s \in \text{prefixes}\,(P) \\ s'' & \text{if } s = s's'' \text{ for some } s' \in P. \end{cases} \tag{20}$$

Let P_X be a balanced X-partition of C, and let P_Y be a balanced Y-partition of C. Define the *X, Y-load balancing map*

$$\mathsf{Balance}_{X,Y} : \mathbf{D}(X) \times \mathbf{D}(Y) \times \mathbf{D}(W) \times \mathbf{D}(T) \to \mathbf{D}(X') \times \mathbf{D}(Y') \times \mathbf{D}(W) \times \mathbf{D}(T) \times \mathbf{D}(Y'') \times \mathbf{D}(X''),$$

by setting

$$\mathsf{Balance}_{X,Y}(\langle x, y, w, t\rangle) := \langle x'(P_X), y'(P_Y), w, t, y''(P_Y), x''(P_X)\rangle.$$

Let $\mathsf{Balance}_{X,Y}(C) = \left\{\mathsf{Balance}_{X,Y}(\mathbf{b}) \mid \mathbf{b} \in C\right\}$. Our algorithm is to run Tetris-Preloaded with $\mathsf{Balance}_{X,Y}(C)$ as the input. Note that the problem has now been "lifted up" to become a 6-dimensional problem!

To analyze the algorithm, for each $x \in P_X$, $y \in P_Y$, define the following sets:

$$\begin{aligned}
\mathcal{F}[x] &= \left\{\langle x, y', w, t, y'', x''\rangle \in \mathsf{Balance}_{X,Y}(C) \mid x'' \neq \lambda\right\} \\
\mathcal{F}[x,y] &= \left\{\langle x', y, w, t, y'', \lambda\rangle \in \mathsf{Balance}_{X,Y}(C) \mid x' \in \text{prefixes}(x) \text{ and } y'' \neq \lambda\right\} \\
C[x,y] &= \left\{\langle x', y', w, t, \lambda, \lambda\rangle \in \mathsf{Balance}_{X,Y}(C) \mid x' \in \text{prefixes}(x) \text{ and } y' \in \text{prefixes}(y)\right\}
\end{aligned}$$

Note that $x' \in \text{prefixes}(x)$ as binary strings means $x' \supseteq x$ when viewed as dyadic intervals. The reader should keep in mind that we use these two notations interchangeably, for sometimes one notation is more succinct than the other. Note also that,

$$\begin{aligned}
|\mathcal{F}[x]| &\leqslant \sqrt{|C|} \\
|\mathcal{F}[x,y]| &\leqslant \sqrt{|C|} \\
|C[x,y]| &\leqslant |C|.
\end{aligned}$$

We will use these facts to bound the number of witnesses of various kinds. To do so, we need a new notation.

Definition F.6 (Conditional witness sets)**.** Let A be an arbitrary attribute, and S be a set of attributes that does not contain A. Then, we use $\mathcal{W}(A \mid \text{some condition on } S)$ to denote the set of A-witnesses $\mathbf{w}$ such that $\pi_S(\mathbf{w})$ satisfies the condition on S. For example, let X, Y, and W be three different attributes, then

$$\begin{aligned}
\mathcal{W}(X \mid Y = y) &= \{\mathbf{w} \mid \mathbf{w} \text{ is an } X\text{-witness with } \pi_Y(\mathbf{w}) = y\} \\
\mathcal{W}(X \mid Y \supseteq y \wedge W = w) &= \{\mathbf{w} \mid \mathbf{w} \text{ is an } X\text{-witness with } \pi_Y(\mathbf{w}) \supseteq y \wedge \pi_W(\mathbf{w}) = w\}.
\end{aligned}$$

When the condition part is empty, $\mathcal{W}(A)$ denotes the total set of A-witnesses.

We are now ready to bound the number of A-witnesses for $A \in \{X'', Y'', W, T, X', Y'\}$.

- The number of X''-witnesses is easy to bound. By definition, an X''-witness $\mathbf{w}$ is a box that is involved in a gap box resolution on X''; this means $\pi_{X'}(\mathbf{w}) \in P_X$. Moreover, X''-witnesses that do not share the X'-value will not resolve with one another to form another X''-witness, because the strings $x \in P_X$ are prefix-free. Consequently,

  $$\left|\mathcal{W}(X'')\right| = \sum_{x \in P_X} \left|\mathcal{W}(X'' \mid X' = x)\right|.$$

  We bound $\left|\mathcal{W}(X'' \mid X' = x)\right|$ for each $x \in P_X$ by proving the following claim.

  Claim 1. For any $x \in P_X$, an X''-witness $\mathbf{w} \in \mathcal{W}(X'' \mid X' = x)$ is supported on $\{Y', Y'', X''\}$, $\{W, X''\}$, and $\{T, X''\}$ by three boxes from $\mathcal{F}[x]$.

  From the claim, the total number of X''-witnesses is at most

  $$\sum_{x \in P_X} \tilde{O}(|\mathcal{F}[x]|^3) = \tilde{O}(|P_X| \cdot \sqrt{|C|}^3) = \tilde{O}(|C|^2).$$

  Claim 1 is easily shown by induction, whose proof we omit.

- Bounding the number of Y''-witnesses requires a new simple but subtle idea. We use the following obvious bound,

  $$\left|\mathcal{W}(Y'')\right| \leqslant \sum_{x \in P_X} \left|\mathcal{W}(Y'' \mid X' \supseteq x)\right|,$$

  and then bound each term on the right hand side by proving the second claim.

Claim 2. For any $x \in P_X$, a witness $\mathbf{w} \in \mathcal{W}(Y'' \mid X' \supseteq x)$ must satisfy the following two conditions

(i) $\mathbf{w}$ either is supported on $\{Y', Y''\}$ and $\{W\}$ by two boxes from $\mathcal{F}[x]$, or is supported on $\{Y', Y'', W\}$ by a box from $\mathcal{F}[x, \pi_{Y'}(\mathbf{w})]$. (Notice that by definition of Balance, for each $\mathbf{w} \in \mathcal{W}(Y'')$, $\pi_{Y'}(\mathbf{w}) \in P_Y$.)

(ii) $\mathbf{w}$ either is supported on $\{Y', Y''\}$ and $\{T\}$ by two boxes from $\mathcal{F}[x]$, or is supported on $\{Y', Y'', T\}$ by a box from $\mathcal{F}[x, \pi_{Y'}(\mathbf{w})]$.

Assuming Claim 2, then for each witness $\mathbf{w} \in \mathcal{W}(Y'' \mid X' \supseteq x)$, the box $\pi_{\{Y',W,T,Y''\}}(\mathbf{w})$ either is (1) supported by three boxes from $\mathcal{F}(x)$, or is (2) supported by two boxes from $\mathcal{F}[x] \cup \mathcal{F}[x, y]$ for some $y \in P_Y$. The number of witnesses satisfying (1) is at most $\tilde{O}(|\mathcal{F}[x]|^3)$. The number of witnesses satisfying (2) is at most $\tilde{O}(1) \sum_{y \in P_Y} (|\mathcal{F}[x]| + |\mathcal{F}[x, y]|)^2$. Consequently,

$$|\mathcal{W}(Y'')| \leqslant \tilde{O}(1) \sum_{x \in P_X} \left(|\mathcal{F}[x]|^3 + \sum_{y \in P_Y} (|\mathcal{F}[x]| + |\mathcal{F}[x, y]|)^2 \right) = \tilde{O}(|C|^2).$$

We next prove Claim 2 above by induction. Consider a witness $\mathbf{w} \in \mathcal{W}(Y'' \mid X' \supseteq x)$. In the base case, $\mathbf{w}$ is *not* the result of a resolution on Y''. In this case, either $\mathbf{w} \in \mathcal{W}(X'' \mid X' = x)$, or $\mathbf{w} \in \mathcal{F}[x, \pi_{Y'}(\mathbf{w})]$. In both cases, by noting Claim 1, it is easy to see that both (i) and (ii) hold.

Next, suppose $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$, where $\mathbf{w}_1$ and $\mathbf{w}_2$ satisfy (i) and (ii). Because the two witnesses resolve on Y'', $\pi_{Y''}(\mathbf{w}_1) \neq \lambda$ and $\pi_{Y''}(\mathbf{w}_2) \neq \lambda$. Thus, $\pi_{Y'}(\mathbf{w}_1) \in P_Y$ and $\pi_{Y'}(\mathbf{w}_2) \in P_Y$. Because the strings in P_Y are prefix-free, we conclude that $\pi_{Y'}(\mathbf{w}_1) = \pi_{Y'}(\mathbf{w}_2) = \pi_{Y'}(\mathbf{w})$. By definition of resolution, either $\pi_{\{Y',Y'',W\}}(\mathbf{w}_1) \subseteq \pi_{\{Y',Y'',W\}}(\mathbf{w})$ or $\pi_{\{Y',Y'',W\}}(\mathbf{w}_2) \subseteq \pi_{\{Y',Y'',W\}}(\mathbf{w})$, and the same holds for $\{Y', Y'', T\}$. The inductive step follows.

- The number of T-witnesses can be bounded by

$$|\mathcal{W}(T)| \leqslant \sum_{x \in P_X} \sum_{y \in P_Y} |\mathcal{W}(T \mid X' \supseteq x, Y' \supseteq y)|.$$

We bound the terms $|\mathcal{W}(T \mid X' \supseteq x, Y' \supseteq y)|$ individually using the following claim.

Claim 3. For a fixed $x \in P_X$ and $y \in P_Y$, a witness $\mathbf{w} \in \mathcal{W}(T \mid X' \supseteq x, Y' \supseteq y)$ must either be (1) supported on $\{W\}$ and $\{T\}$ by two boxes from $\mathcal{F}[x] \cup \mathcal{F}[x, y]$, or be (2) supported on $\{W, T\}$ by a box from $C[x, y]$.

Assuming the claim holds, then the number of T-witnesses is bounded by

$$\begin{aligned} |\mathcal{W}(T)| &\leqslant \sum_{x \in P_X, y \in P_Y} |\mathcal{W}(T \mid X' \supseteq x, Y' \supseteq y)| \\ &\leqslant \sum_{x \in P_X, y \in P_Y} \tilde{O}\left((|\mathcal{F}[x]| + |\mathcal{F}[x, y]|)^2 + |C[x, y]| \right) \\ &\leqslant \sum_{x \in P_X, y \in P_Y} \tilde{O}(|\sqrt{|C|}|^2 + |C|) = \tilde{O}(|C|^2). \end{aligned}$$

We next prove Claim 3. For a fixed $x \in P_X$ and $y \in P_Y$, consider a witness $\mathbf{w} \in \mathcal{W}(T \mid X' \supseteq x, Y' \supseteq y)$. Since our resolutions are ordered, $\pi_{X''}(\mathbf{w}) = \pi_{Y''}(\mathbf{w}) = \lambda$. In the base case, $\mathbf{w}$ could be either an input gap box or a result of a gap box resolution on either X'' or Y''. If it is an input gap box, then $\mathbf{w} \in C[x, y]$. If it is a result of gap box resolution on X'', then $\pi_{X'}(\mathbf{w}) = x$ and Claim 3 holds because Claim 1 holds. If it is a result of a gap box resolution on Y'', then $\pi_{Y'}(\mathbf{w}) = y$ and the claim holds because Claim 2 holds. The inductive step is very similar to that of Claim 2. Suppose that Claim 3 holds for $\mathbf{w}_1$ and $\mathbf{w}_2$. Either $\pi_{\{W,T\}}(\mathbf{w}_1) \subseteq \pi_{\{W,T\}}(\mathbf{w})$ or $\pi_{\{W,T\}}(\mathbf{w}_2) \subseteq \pi_{\{W,T\}}(\mathbf{w})$, and hence the claim holds for $\mathbf{w}$.

- Finally, we bound the number of gap box resolutions on W, X', and Y'. Each witness $\mathbf{w}$ is supported by some box in C on each one of the attributes W, X' and Y'. While W can take $\tilde{O}(|C|)$ values, each one of X' and Y' takes only $\tilde{O}(\sqrt{|C|})$ values. The total number of W, Y' and X'-witnesses is bounded by $\tilde{O}(|C|^2)$, as desired.

F.5 Tetris-Preloaded with load balancing

In this section, we generalize the ideas presented in the previous section to analyze Tetris-Preloaded-LB with the load balancing map idea incorporated. Algorithm 5 has the key steps. The load balancing map is computed as a preprocessing step before calling Tetris-Preloaded on the mapped boxes, which have been "lifted up" to a higher-dimensional space.

For $i \in [n-2]$, let P_i denote the balanced A_i-partition used in the algorithm. Define the load-balancing map

$$\mathsf{Balance}_{A_1,\ldots,A_{n-2}} : \mathbf{D}(A_1) \times \cdots \times \mathbf{D}(A_n) \to \mathbf{D}(A'_1) \times \cdots \times \mathbf{D}(A'_{n-2}) \times \mathbf{D}(A_n) \times \mathbf{D}(A_{n-1}) \times \mathbf{D}(A''_{n-2}) \times \cdots \times \mathbf{D}(A''_1),$$

by setting

$$\mathsf{Balance}_{A_1,\ldots,A_{n-2}}(\langle b_1,\ldots,b_n\rangle) := \langle b'_1(P_1),\ldots,b'_{n-2}(P_{n-2}), b_n, b_{n-1}, b''_{n-2}(P_{n-2}),\ldots,b''_1(P_1)\rangle.$$

(Recall the definitions of the functions $s'(P)$ and $s''(P)$ from (19) and (20).) Also, naturally define

$$\mathsf{Balance}_{A_1,\ldots,A_{n-2}}(C) = \left\{\mathsf{Balance}_{A_1,\ldots,A_{n-2}}(\mathbf{b}) \mid \mathbf{b} \in C\right\}.$$

Algorithm 5 Tetris-Preloaded-LB, (i.e. Tetris-Preloaded with load balancing)

Input: A set of boxes C in the domain $\mathbf{D}(A_1) \times \cdots \times \mathbf{D}(A_n)$
Output: Output tuples for the BCP on C

1: **For** $i \leftarrow 1$ to $n-2$ **do** ▷ $\tilde{O}(|C|)$-time
2: $P_i \leftarrow$ a balanced A_i-partition
3: $\sigma \leftarrow (A'_1, A'_2, \ldots, A'_{n-2}, A_n, A_{n-1}, A''_{n-2}, A''_{n-3}, \ldots, A''_1)$
4: $\mathcal{B} \leftarrow \mathsf{Balance}_{A_1,\ldots,A_{n-2}}(C)$ ▷ $\tilde{O}(|C|)$-time
5: **Return** Tetris-Preloaded($\mathcal{B}$) ▷ SAO $= \sigma$

The main theorem of this section is the following.

Theorem F.7 (Tetris-Preloaded-LB runs in time $\tilde{O}(|C|^{n/2}+Z)$)**.** *For any set of boxes C in n dimensions,* Tetris-Preloaded-LB *solves* BCP *on input C in time $\tilde{O}(|C|^{n/2}+Z)$, where Z is the output size.*

Proof. For $i \in [n-2]$, let P_i denote the balanced A_i-partition as computed in line 2 of the algorithm. For $m = 1$ to $n-2$, define the following sets for every tuple $(a_1,\ldots,a_m) \in P_1 \times \cdots \times P_m$:

$$\begin{aligned}
\mathcal{F}[a_1,\ldots,a_m] &:= \left\{\mathbf{b} \in \mathsf{Balance}_{A_1,\ldots,A_{n-2}}(C) \;\middle|\; \pi_{A'_i}(\mathbf{b}) \supseteq a_i \wedge \pi_{A''_i}(\mathbf{b}) = \lambda, \forall i \in [m-1] \wedge \pi_{A'_m}(\mathbf{b}) = a_m \wedge \pi_{A''_m}(\mathbf{b}) \neq \lambda\right\} \\
C[a_1,\ldots,a_m] &:= \left\{\mathbf{b} \in \mathsf{Balance}_{A_1,\ldots,A_{n-2}}(C) \;\middle|\; \pi_{A'_i}(\mathbf{b}) \supseteq a_i \wedge \pi_{A''_i}(\mathbf{b}) = \lambda, \forall i \in [m]\right\}
\end{aligned}$$

For notational convenience, also define the following set:

$$\mathcal{G}[a_1,\ldots,a_m] := \bigcup_{i=1}^{m} \mathcal{F}[a_1,\ldots,a_i].$$

It should be noted that $\mathsf{Balance}_{A_1,\ldots,A_{n-2}}(C) = \bigcup_{(a_1,\ldots,a_{n-2}) \in P_1 \times \cdots \times P_{n-2}} \mathcal{G}[a_1,\ldots,a_{n-2}] \cup C[a_1,\ldots,a_{n-2}]$. Recalling the notation defined in (11), and the definition of balanced dimension partitions, we have, for every $m \in [n-2]$ and every tuple $(a_1,\ldots,a_m) \in P_1 \times \cdots \times P_m$,

$$\begin{aligned}
\left|\mathcal{F}[a_1,\ldots,a_m]\right| &\leqslant \left|C_{\subset a_m}(A_m)\right| \leqslant \sqrt{|C|}, \\
\left|\mathcal{G}[a_1,\ldots,a_m]\right| &\leqslant m\sqrt{|C|} = \tilde{O}(\sqrt{|C|}), \\
\left|C[a_1,\ldots,a_m]\right| &\leqslant |C|.
\end{aligned}$$

In addition, for each $i \in [n]$, define S_i as the following set of attributes:

$$S_i = \begin{cases} \{A'_i, A''_i\} & \text{if } i \in [n-2] \\ \{A_i\} & \text{otherwise.} \end{cases} \tag{21}$$

Recall that the SAO used by Tetris-Preloaded-LB is $(A'_1, A'_2, \ldots, A'_{n-2}, A_n, A_{n-1}, A''_{n-2}, A''_{n-3}, \ldots, A''_1)$. We will bound the total number of gap box resolutions and then use Theorem C.5. Note that when an output point is found and Tetris inserts it back to $\mathcal{A}$ (Line 6 of Algorithm 2), this output point is now in the higher dimension of the Balance map.

Recall the definition of (conditional) witness sets (Definition F.6). We write the total number of gap box resolutions as

$$\sum_{m=1}^{n-2} |\mathcal{W}(A''_m)| + |\mathcal{W}(A_n)| + |\mathcal{W}(A_{n-1})| + \sum_{m=1}^{n-2} |\mathcal{W}(A'_m)|.$$

In what follows, we bound each of the above terms.

Bounding $|\mathcal{W}(A''_m)|$ for $m \in [n-2]$ To bound $|\mathcal{W}(A''_m)|$, we use the following inequality:

$$|\mathcal{W}(A''_m)| \leqslant \sum_{a_1 \in P_1} \cdots \sum_{a_{m-1} \in P_{m-1}} |\mathcal{W}(A''_m \mid a_i \subseteq A'_i, \forall i \in [m-1])|.$$

Notice that in order for a witness $\mathbf{w}$ to be involved in an A''_m-resolution, it must be the case that $\pi_{A''_i}(\mathbf{w}) = \lambda$, for all $i \in [m-1]$. In particular,

$$\mathsf{support}(\mathbf{w}) \subseteq \{A'_1, \ldots, A'_{n-2}, A_n, A_{n-1}, A''_{n-2}, \ldots, A''_m\} = \{A'_1, \ldots, A'_{m-1}\} \cup S_m \cup \cdots \cup S_n.$$

It must also hold that $\pi_{A'_m}(\mathbf{w}) \in P_m$, which is a fact we will use below. The idea is to bound each term $|\mathcal{W}(A''_m \mid a_i \subseteq A'_i, \forall i \in [m-1])|$ individually, by proving the following claims by induction on $m = 1, \ldots, n-2$.

Claim m. For every tuple $(a_1, \ldots, a_{m-1}) \in P_1 \times \cdots \times P_{m-1}$ and for every A''_m-witness

$$\mathbf{w} \in \mathcal{W}(A''_m \mid a_i \subseteq A'_i, \forall i \in [m-1]),$$

the following property holds

- For every $i \in \{m+1, \ldots, n\}$, either $\mathbf{w}$ is supported on S_m and S_i by two boxes in $\mathcal{G}[a_1, \ldots, a_{m-1}]$, or $\mathbf{w}$ is supported on $S_m \cup S_i$ by some box in $\mathcal{F}[a_1, \ldots, a_{m-1}, \pi_{A'_m}(\mathbf{w})]$.

Assuming Claim m is correct, the term $|\mathcal{W}(A''_m \mid a_i \subseteq A'_i, \forall i \in [m-1])|$ is bounded as follows. For each $\mathbf{w} \in \mathcal{W}(A''_m \mid a_i \subseteq A'_i, \forall i \in [m-1])$, define the projection of $\mathbf{w}$ onto the set $\mathsf{support}(\mathbf{w}) - \{A'_1, \ldots, A'_{m-1}\}$:

$$\overline{\mathbf{w}} = \pi_{\mathsf{support}(\mathbf{w}) - \{A'_1, \ldots, A'_{m-1}\}}(\mathbf{w}).$$

Then,

- Either $\overline{\mathbf{w}}$ is supported by an integral cover of size $\leqslant n - m + 1$ using the box set $\mathcal{G}(a_1, \ldots, a_{m-1})$.
- Or $\overline{\mathbf{w}}$ is supported by an integral cover of size $\leqslant n - m$ using the box set $\mathcal{G}(a_1, \ldots, a_{m-1}, \pi_{A'_m}(\mathbf{w}))$.

This fact helps bound the total number of such witnesses by

$$\tilde{O}(1) \sum_{a_1 \in P_1} \cdots \sum_{a_{m-1} \in P_{m-1}} \left[\left(\sqrt{|C|}\right)^{n-m+1} + \sum_{a_m \in P_m} \left(\sqrt{|C|}\right)^{n-m} \right] = \tilde{O}(|C|^{n/2}).$$

We next prove Claim m by induction on gap box resolutions and on m.

In the base case, $\mathbf{w}$ could be either an input gap box or a result of a gap box resolution on A''_l for some $l < m$ (assuming $m > 1$). If $\mathbf{w}$ is an input gap box, then $\mathbf{w} \in \mathcal{F}[a_1, \ldots, a_{m-1}, \pi_{A'_m}(\mathbf{w})]$ and the claim holds. If $m > 1$ and $\mathbf{w}$ is a result of a gap box resolution on A''_l for some $l < m$, then $A'_l(\mathbf{w}) = a_l$. By the induction hypothesis, Claim l holds, which means $\mathbf{w}$ is supported on each one of the sets $S_{l+1}, \ldots, S_n$ by some box in $\mathcal{G}(a_1, \ldots, a_l) \subseteq \mathcal{G}(a_1, \ldots, a_{m-1})$, and hence Claim m holds in this case.

For the inductive step, let's assume that Claim m holds for two witnesses $\mathbf{w}_1$ and $\mathbf{w}_2$, and we will prove that it holds for $\mathbf{w} := \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$. Because the two witnesses resolve on A''_m, we have $\pi_{A''_m}(\mathbf{w}_1) \neq \lambda$ and $\pi_{A''_m}(\mathbf{w}_2) \neq \lambda$. This means $\pi_{A'_m}(\mathbf{w}_1) \in P_m$ and $\pi_{A'_m}(\mathbf{w}_2) \in P_m$. But the strings in P_m are prefix-free, leading to $\pi_{A'_m}(\mathbf{w}_1) = \pi_{A'_m}(\mathbf{w}_2) = \pi_{A'_m}(\mathbf{w})$. For every $i \in \{m+1, \ldots, n\}$, either $\pi_{S_i \cup S_m}(\mathbf{w}_1) \subseteq \pi_{S_i \cup S_m}(\mathbf{w})$ or $\pi_{S_i \cup S_m}(\mathbf{w}_2) \subseteq \pi_{S_i \cup S_m}(\mathbf{w})$ by definition of resolution. The inductive step follows.

Bounding $|\mathcal{W}(A_{n-1})|$ We bound the number of A_{n-1}-witnesses using the following inequality

$$|\mathcal{W}(A_{n-1})| \leqslant \sum_{a_1 \in P_1} \cdots \sum_{a_{n-2} \in P_{n-2}} |\mathcal{W}(A_{n-1} \mid a_i \subseteq A'_i, \forall i \in [n-2])|.$$

Notice that in order for a witness $\mathbf{w}$ to be involved in an A_{n-1}-resolution, it must be the case that $\pi_{A''_i}(\mathbf{w}) = \lambda$, for all $i \in [n-2]$. So, we bound $|\mathcal{W}(A_{n-1})|$ by bounding the terms $|\mathcal{W}(A_{n-1} \mid a_i \subseteq A'_i, \forall i \in [n-2])|$, assisted by the following claim.

Claim $n-1$. For every tuple $(a_1, \ldots, a_{n-2}) \in P_1 \times \cdots \times P_{n-2}$ and for every A_{n-1}-witness

$$\mathbf{w} \in \mathcal{W}(A_{n-1} \mid a_i \subseteq A'_i, \forall i \in [n-2]),$$

the following property holds:

- Either $\mathbf{w}$ is supported on $\{A_n\}$ and $\{A_{n-1}\}$ by two boxes in $\mathcal{G}[a_1, \ldots, a_{n-2}]$, or $\mathbf{w}$ is supported on $\{A_n, A_{n-1}\}$ by some box in $C[a_1, \ldots, a_{n-2}]$.

Assuming Claim $n-1$ is correct, the number of A_{n-1}-witnesses is bounded by

$$\tilde{O}(1) \sum_{a_1 \in P_1} \cdots \sum_{a_{n-2} \in P_{n-2}} \left[\sqrt{|C|}^2 + |C| \right] = \tilde{O}(|C|^{n/2}).$$

Claim $n-1$ can be proved in a very similar way to Claim m where $m = n-1$ and $C[a_1, \ldots, a_{n-2}]$ replaces $\mathcal{F}[a_1, \ldots, a_{m-1}, \pi_{A'_m}(\mathbf{w})]$.

Bounding $|\mathcal{W}(A'_m)|$ for $m \in [n-2]$ and $|\mathcal{W}(A_n)|$ Finally, we will bound the total number of $A'_1, \ldots, A'_{n-2}$ and A_n-witnesses. Every witness $\mathbf{w}$ that is involved in a gap box resolution must be supported on each attribute by some box in C. Attribute A_n can take up to $\tilde{O}(|C|)$ different values. However, from the definition of balanced domain partition, each one of the attributes $\{A'_1, \ldots, A'_{n-2}\}$ takes only $\tilde{O}(\sqrt{|C|})$ values. Therefore, the total number of $A'_1, \ldots, A'_{n-2}$ and A_n-witnesses is bounded by $\tilde{O}(|C|^{n/2})$, as desired. □

Corollary F.8 (Subsumed by [13, 56])**.** *Given a set of boxes $\mathcal{B}$ in n dimensions, Klee's measure problem over the Boolean semiring can be solved in time $\tilde{O}(|\mathcal{B}|^{n/2})$.*

F.6 Tetris-Reloaded with load balancing

In this section, we generalize the ideas presented in the previous section to analyze Tetris-Reloaded with the load balancing map idea incorporated. The main theorem we will show is the following.

Theorem F.9 (Tetris-Reloaded-LB runs in time $\tilde{O}(|C_\square|^{n/2} + Z)$)**.** *For any set of boxes $\mathcal{B}$ in n dimensions,* Tetris-Reloaded-LB *solves* BCP *on input $\mathcal{B}$ in time $\tilde{O}(|C_\square|^{n/2} + Z)$. Here, $C_\square$ is any optimal box certificate for the instance, and Z is the output size.*

Corollary F.10. Tetris-Reloaded-LB *evaluates any join query Q over n attributes in time $\tilde{O}(|C_\square|^{n/2} + Z)$, where $C_\square$ is an optimal box certificate for the join instance, and Z is the output size.*

In Tetris-Reloaded with load balancing (Algorithm 6), the dimension partitions are done with respect to the set C of input gap boxes that have actually been loaded from $\mathcal{B}$. Once new boxes are loaded from $\mathcal{B}$ into C, the partitions are updated as follows: If an interval x of any partition P_i becomes *heavy* (i.e., if $|C_{\subset x}(A_i)|$ exceeds $\sqrt{|C|}$), x gets replaced by $x0$ and $x1$. Later on when new boxes that are not contained in the x-layer are loaded into C, $|C|$ is going to increase while $|C_{\subset x}(A_i)|$ remains the same. However, even if $\sqrt{|C|}$ exceeds $|C_{\subset x}(A_i)|$ again, we do NOT return x into P_i instead of $x0$ and $x1$. In other words, we only allow partitions P_i to "expand". Lemma F.11 below shows that the partition sizes are still going to remain within $\tilde{O}(\sqrt{|C|})$. The proof of Theorem F.9 is going to rely on the fact that partitions expand only.

Whenever any partition P_i is updated, boxes in the partitioned space must be updated accordingly to reflect the new A_i-partition. In particular, whenever some x in P_i is replaced by $x0$ and $x1$, the knowledge base $\mathcal{A}$ must be updated by invoking $\mathsf{Update\text{-}Balance}_{A_i,x}(\mathcal{A})$ (in Line 17 of Algorithm 6). $\mathsf{Update\text{-}Balance}_{A_i,x}(\mathcal{A})$ works as follows:

- For every $\mathbf{a} \in \mathcal{A}$ such that $\pi_{A_i'}(\mathbf{a}) = x$ and $\pi_{A_i''}(\mathbf{a}) = by$ from some $b \in \{0,1\}, y \in \{0,1\}^*$, $\pi_{A_i'}(\mathbf{a})$ becomes xb and $\pi_{A_i''}(\mathbf{a})$ becomes y.

$\mathsf{Update\text{-}Balance}_{A_i,x}(\mathcal{A})$ does not update any box $\mathbf{a} \in \mathcal{A}$ for which $\pi_{A_i'}(\mathbf{a}) \neq x$ or $\pi_{A_i''}(\mathbf{a}) = \lambda$. Notice that because partitions expand only, no box $\mathbf{a} \in \mathcal{A}$ can be updated more than $\tilde{O}(1)$ times.

Algorithm 6 Tetris-Reloaded-LB, (i.e. Tetris-Reloaded with Load Balancing)

Input: Oracle access to a set of boxes $\mathcal{B}$ in the domain $\mathbf{D}(A_1) \times \cdots \times \mathbf{D}(A_n)$
Output: Output tuples for the BCP on $\mathcal{B}$

1: $C \leftarrow \varnothing$ ▷ C is a set of boxes in the domain $\mathbf{D}(A_1) \times \cdots \times \mathbf{D}(A_n)$.
2: **For** $i \leftarrow 1$ to $n-2$ **do**
3: $P_i \leftarrow \{\lambda\}$
4: $\sigma \leftarrow (A_1', A_2', \ldots, A_{n-2}', A_n, A_{n-1}, A_{n-2}'', A_{n-3}'', \ldots, A_1'')$
5: $\mathcal{A} \leftarrow \varnothing$ ▷ $\mathcal{A}$ is a set of boxes in the domain $\mathbf{D}(A_1') \times \cdots \times \mathbf{D}(A_{n-2}') \times \mathbf{D}(A_n) \times \mathbf{D}(A_{n-1}) \times \mathbf{D}(A_{n-2}'') \times \cdots \times \mathbf{D}(A_1'')$.
6: $(v, \mathbf{w}) \leftarrow \mathsf{TetrisSkeleton}(\langle \lambda, \ldots, \lambda \rangle)$ ▷ SAO $= \sigma$
7: **While** $v =$ FALSE **do**
8: $\mathcal{B}' \leftarrow \{\mathbf{b} \in \mathcal{B} \mid \mathsf{Balance}_{A_1,\ldots,A_{n-2}}(\mathbf{b}) \supseteq \mathbf{w}\}$
9: **If** $\mathcal{B}' = \varnothing$ **then**
10: **Report** $\mathbf{w}$ as an output tuple
11: $\mathcal{B}' \leftarrow \{\mathbf{w}\}$
12: **else**
13: $C \leftarrow C \cup \mathcal{B}'$
14: **For** $i \leftarrow 1$ to $n-2$ **do**
15: **While** $\exists x \in P_i \mid |C_{\subset x}(A_i)| > \sqrt{|C|}$ **do**
16: $P_i \leftarrow P_i - \{x\} \cup \{x0, x1\}$ ▷ Update the partition P_i
17: $\mathsf{Update\text{-}Balance}_{A_i,x}(\mathcal{A})$ ▷ Re-balance $\mathcal{A}$ with the new P_i
18: $\mathcal{A} \leftarrow \mathcal{A} \cup \mathsf{Balance}_{A_1,\ldots,A_{n-2}}(\mathcal{B}')$
19: $(v, \mathbf{w}) \leftarrow \mathsf{TetrisSkeleton}(\langle \lambda, \ldots, \lambda \rangle)$ ▷ SAO $= \sigma$

Lemma F.11. *At the end of execution of* Tetris-Reloaded-LB, $|P_i| = \tilde{O}(\sqrt{|C|})$ *for each* $i \in [n-2]$.

Proof. Consider an arbitrary but fixed $i \in [n-2]$. When Tetris-Reloaded-LB starts, $|P_i| = 1$. Whenever $|P_i|$ increases by one (in Line 16 of Algorithm 6), there must be some string $x \in \{0,1\}^*$ such that $|C_{\subset x}(A_i)| > \sqrt{|C|}$. Consider an arbitrary but fixed integer $k > 0$. From the moment $|C|$ reaches k until the moment it reaches $4k$, the increase in $|P_i|$ can be bounded as follows: Whenever $|P_i|$ increases by one, there must be some string x that satisfies $|C_{\subset x}(A_i)| > \sqrt{k}$.

However, for every box $\mathbf{c} \in C$ (whose size is $\leqslant 4k$), there are $\tilde{O}(1)$ strings x that satisfy $\mathbf{c} \in C_{\sqsubset x}(A_i)$. Hence, the increase in $|P_i|$ is bounded by $\tilde{O}(\sqrt{k})$.

At the end of execution of Tetris-Reloaded-LB, the total size of P_i is bounded by the sum of $\tilde{O}(\sqrt{k})$ over all values $k \in \left\{4^0, 4^1, \ldots, 4^{1/2\lceil \log_2 |C| \rceil}\right\}$:

$$|P_i| \leqslant \tilde{O}(1)\left[2^0 + 2^1 + \cdots + 2^{1/2\lceil \log_2 |C| \rceil}\right] \leqslant \tilde{O}(1) 2^{1/2\lceil \log_2 |C| \rceil}\left[1 + 1/2 + 1/4 + \cdots\right] \leqslant \tilde{O}(\sqrt{|C|}).$$

□

Proof of Theorem F.9. Because Tetris performs ordered resolutions only, a box cannot be involved in a resolution on some attribute unless all subsequent attributes in the SAO are λ's. If some box $\mathbf{w}$ has a λ in attribute A_i'', then $\pi_{A_i''}(\mathbf{w})$ is going to remain λ even after the A_i-partition P_i is updated. This is because partitions expand only: any $x \in P_i$ can be replaced by $x0$ and $x1$, but $x0$ and $x1$ cannot be replaced back by x. As a result, resolutions remain ordered throughout the whole execution of Tetris-Reloaded-LB.

During the execution of Tetris-Reloaded-LB (Algorithm 6), a light interval $x \in P_i$ might get heavy (line 15), in which case x has to be split (line 16). Given $i \in [n-2]$ and a binary string $x \in \{0,1\}^*$, let $\mathcal{U}_i(x)$ refer to the last value of $C_{\sqsubset x}(A_i)$ while x was still a light interval (i.e., the last $C_{\sqsubset x}(A_i)$ while $|C_{\sqsubset x}(A_i)|$ has not yet exceeded $\sqrt{|C|}$ at any previous step). By definition, $\left|\mathcal{U}_i(x)\right| \leqslant \sqrt{|C|}$. Define $\mathcal{V}_i(x)$ as follows:

$$\mathcal{V}_i(x) := \bigcup_{y \in \text{prefixes}(x)} \mathcal{U}_i(y)$$

By definition, $\left|\mathcal{V}_i(x)\right| = \tilde{O}(\sqrt{|C|})$.

Given a variable X, we will use $\hat{X}$ to refer to the *final value* of X at the end of execution of Tetris-Reloaded-LB. In particular, we will be using $\hat{P}_1, \ldots, \hat{P}_{n-2}$, and $\hat{C}$ to refer to the final values of $P_1, \ldots, P_{n-2}$, and C respectively. Given a box $\mathbf{w} \in \mathcal{A}$, invocations of $\textsf{Update-Balance}_{A_i,x}(\mathcal{A})$ might change $\pi_{A_i'}(\mathbf{w})$ and $\pi_{A_i''}(\mathbf{w})$. However, we will be using $\hat{\mathbf{w}}$ to refer to the final value of $\mathbf{w}$ (i.e. after the last invocation of Update-Balance). Similarly, we will be using $\hat{\textsf{Balance}}_{A_1,\ldots,A_{n-2}}$ to refer to the final load balancing map (i.e. with respect to the final partitions $\hat{P}_1, \ldots, \hat{P}_{n-2}$):

$$\hat{\textsf{Balance}}_{A_1,\ldots,A_{n-2}}(\langle b_1, \ldots, b_n \rangle) := \langle b_1'(\hat{P}_1), \ldots, b_{n-2}'(\hat{P}_{n-2}), b_n, b_{n-1}, b_{n-2}''(\hat{P}_{n-2}), \ldots, b_1''(\hat{P}_1) \rangle.$$

$$\hat{\textsf{Balance}}_{A_1,\ldots,A_{n-2}}(C) = \left\{\hat{\textsf{Balance}}_{A_1,\ldots,A_{n-2}}(\mathbf{b}) \mid \mathbf{b} \in C\right\}.$$

By Lemma F.11, $\left|\hat{P}_i\right| = \tilde{O}(\sqrt{\left|\hat{C}\right|})$ for all $i \in [n-2]$. For $m = 1$ to $n-2$, define the following sets for every tuple $(a_1, \ldots, a_m) \in \hat{P}_1 \times \cdots \times \hat{P}_m$:

$$\begin{aligned}
\mathcal{F}[a_1, \ldots, a_m] &:= \left\{\mathbf{b} \in \hat{\textsf{Balance}}_{A_1,\ldots,A_{n-2}}(\mathcal{V}_m(a_m)) \quad \middle| \quad \pi_{A_i'}(\mathbf{b}) \supseteq a_i \wedge \pi_{A_i''}(\mathbf{b}) = \lambda, \forall i \in [m-1]\right\} \\
C[a_1, \ldots, a_m] &:= \left\{\mathbf{b} \in \hat{\textsf{Balance}}_{A_1,\ldots,A_{n-2}}(C) \quad \middle| \quad \pi_{A_i'}(\mathbf{b}) \supseteq a_i \wedge \pi_{A_i''}(\mathbf{b}) = \lambda, \forall i \in [m]\right\}
\end{aligned}$$

For notational convenience, also define the following set:

$$\mathcal{G}[a_1, \ldots, a_m] := \bigcup_{i=1}^{m} \mathcal{F}[a_1, \ldots, a_i].$$

Notice that $\left|\mathcal{F}(a_1, \ldots, a_m)\right| \leqslant \left|\mathcal{V}_m(a_m)\right| = \tilde{O}(\sqrt{\left|\hat{C}\right|})$, $\left|\mathcal{G}(a_1, \ldots, a_m)\right| = \tilde{O}(\sqrt{\left|\hat{C}\right|})$, and $\left|C(a_1, \ldots, a_m)\right| \leqslant \left|\hat{C}\right|$.

In addition, for each $i \in [n]$, define S_i as the following set of attributes:

$$S_i = \begin{cases} \{A_i', A_i''\} & \text{if } i \in [n-2] \\ \{A_i\} & \text{otherwise.} \end{cases}$$

We prove the following claim by induction on $m = 1, \ldots, n-2$.

Claim m**.** For every tuple $(a_1, \ldots, a_{m-1}) \in \hat{P}_1 \times \cdots \times \hat{P}_{m-1}$ and for every A''_m-witness $\mathbf{w}$ that satisfies

$$a_i \subseteq \pi_{A'_i}(\hat{\mathbf{w}}) \text{ for each } i \in [m-1],$$

the following property holds:

- For every $i \in \{m+1, \ldots, n\}$, either $\hat{\mathbf{w}}$ is supported on S_m and S_i by two boxes in $\mathcal{G}[a_1, \ldots, a_{m-1}]$, or $\hat{\mathbf{w}}$ is supported on $S_m \cup S_i$ by some box in $\mathcal{F}[a_1, \ldots, a_{m-1}, a_m]$ for every $a_m \in \hat{P}_m$ that satisfies $a_m \subseteq \pi_{A'_m}(\hat{\mathbf{w}})$.

The rest of the proof is similar to that of Theorem 4.9. □

Corollary F.12. *Given a set of boxes $\mathcal{B}$ in n dimensions, Klee's measure problem over the Boolean semiring can be solved in time $\tilde{O}(|C_\square|^{n/2})$, where $C_\square$ is any optimal box certificate for $\mathcal{B}$.*

By Definition 3.4, we always have $C_\square \subseteq \mathcal{B}$ and hence $|C_\square| \leqslant |\mathcal{B}|$. The following proposition shows that $|C_\square|$ can be unboundedly smaller than $|\mathcal{B}|$.

Proposition F.13. (There is a class of input instances for which $|C_\square|$ is unboundedly smaller than $|\mathcal{B}|$). *For every integer $n \geqslant 1$ and for every integer $b > 0$, there exists a set $\mathcal{B}$ of n-dimensional boxes such that $|\mathcal{B}| = b$ and $|C_\square| = 1$.*

Proof. As long as $\mathcal{B}$ contains $\langle \lambda, \ldots, \lambda \rangle$, $C_\square = \{\langle \lambda, \ldots, \lambda \rangle\}$. □

F.7 Achieving a runtime of $\tilde{O}(|C_\square|^{\frac{B+1}{2}} + Z)$ for queries with maximum block size B

Given a hypergraph, an *articulation point* (alternatively, a *cut vertex*) is a vertex whose removal increases the number of connected components. A *biconnected component* is a component that does not contain any articulation points. A maximal biconnected component is called a *block*. The blocks of any connected hypergraph can be arranged into a tree called the *block tree*.

Let $\mathcal{B}$ be a BCP instance whose supporting hypergraph $\mathcal{H}(\mathcal{B})$ has a maximum block size of B. We define a load balancing map described by a GAO σ for $\mathcal{B}$ as follows. Initialize σ to be empty (i.e., $\sigma \leftarrow ()$). Construct a block tree for $\mathcal{H}(\mathcal{B})$ and repeat the following two steps until the block tree is empty.

- Let L be a leaf block in the block tree. L can maximally share one attribute with other blocks. Let $A_1, \ldots, A_k$ be the attributes that belong to L and only L. For $i \in [k-2]$, let P_i denote a balanced A_i-partition and let A'_i, A''_i be the two attributes that replace A_i in the load balancing map. Update σ as follows:

$$\sigma \leftarrow (A'_1, A'_2, \ldots, A'_{k-2}, A_k, A_{k-1}, A''_{k-2}, A''_{k-3}, \ldots, A''_1, \sigma)$$

- Remove L from the block tree.

Notice that if L above shares one attribute with other blocks, then $k \leqslant B-1$. Every gap box resolvent $\mathbf{w}$ is supported on the shared attribute by some input gab box, while the projection of $\mathbf{w}$ on the k non-shared attributes can have $\tilde{O}(|C_\square|^{k/2})$ combinations (thanks to load balancing). The total runtime would be $\tilde{O}(|C_\square|^{\frac{k}{2}+1} + Z)$. Alternatively if L does not share any attribute (i.e. L is the root of the block tree), then $k \leqslant B$ and load balancing takes $\tilde{O}(|C_\square|^{\frac{k}{2}} + Z)$. The next corollary follows.

Corollary F.14. *Let $\mathcal{B}$ be a* BCP *instance, B the maximum block size of $\mathcal{H}(\mathcal{B})$, $C_\square$ be an optimal box certificate for $\mathcal{B}$, and Z the output size. If* Tetris-Reloaded-LB *uses the load balancing map and the attribute order σ described above, it runs in time $\tilde{O}(|C_\square|^{\frac{B+1}{2}} + Z)$.*

The above result can obviously be specialized to join queries.

F.8 An example of an unbalanced certificate

We present an example of an unbalanced certificate C for $n = 3$. (For the definition of *balanced*, check Section 4.5.1.) In fact, the boxes of C will correspond to gap boxes for the triangle query $R(A,B) \bowtie S(B,C) \bowtie T(A,C)$. Fix a parameter N. We will use O (and E) to denote the set of odd (even resp.) values in $[N]$. Let C be the union of the following three sets:

$$C_1 = \{\langle o_1, o_2, \lambda\rangle \mid o_1, o_2 \in O\} \cup \{\langle e_1, e_2, \lambda\rangle \mid e_1, e_2 \in E\},$$
$$C_2 = \{\langle \lambda, o_1, o_2\rangle \mid o_1, o_2 \in O\} \cup \{\langle \lambda, e_1, e_2\rangle \mid e_1, e_2 \in E\},$$
$$C_3 = \{\langle o_1, \lambda, o_2\rangle \mid o_1, o_2 \in O\} \cup \{\langle e_1, \lambda, e_2\rangle \mid e_1, e_2 \in E\}.$$

It is not hard to see that the boxes of $C = C_1 \cup C_2 \cup C_3$ cover the entire space, and that every box of C is necessary to cover some point in the space.

We cannot partition the domain of A since the $\Theta(|C|)$ boxes of C_2 span that entire domain. At the same time, we cannot take the whole domain of A as a single part because that part would contain $\Theta(|C|)$ boxes. The example is symmetric so that the same problem occurs when we consider B or C.

G Omitted details from Section 5

In this section, we present the proofs of the lower bounds that were claimed without proof in Section 5.

G.1 The $\Omega(N^{n/2} + Z)$ lowerbound for Tree Ordered Geometric Resolution

We show that Tree Ordered Geometric Resolution is not powerful enough to recover Yannakakis' result on acyclic queries (let alone Theorem 4.4).

Theorem G.1 (There is a class of acyclic input instances with $w = 1$ on which every Tree Ordered Geometric Resolution algorithm runs in time $\Omega(N^{n/2} + Z)$**).** *For every integer $n > 1$ and for every integer $c > 0$, there exists an acyclic join query Q with n attributes, size $N = \Theta(c)$ and of treewidth of 1, such that Q satisfies the following conditions:*

- *The output of Q is empty.*
- *For every GYO elimination order σ of Q that induces treewidth of 1, every* Tree Ordered Geometric Resolution *proof (with ordering σ) of $\langle \lambda, \ldots, \lambda\rangle$ consists of $\Omega(N^{n/2})$ resolutions.*

Proof. Let $\text{vars}(Q)$ be $\{A_1, \ldots, A_n\}$. Let $\text{atoms}(Q)$ be $\{R_1, \ldots, R_{n+1}\}$. Let R_1 be a unary relation over A_1 that contains all odd values, and R_{n+1} be a unary relation over A_n that contains all even values. For each $i \in \{2, \ldots, n\}$, let R_i be a binary relation between A_{i-1} and A_i that contains all pairs of equal parity. More precisely, the corresponding gap boxes are:

$$\mathcal{B}(R_1) := \Big\{\mathbf{c} \quad | \quad \pi_{A_1}(\mathbf{c}) \in \{0,1\}^{d-1}0 \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n] - \{1\}\Big\},$$

and

$$\mathcal{B}(R_{n+1}) := \Big\{\mathbf{c} \quad | \quad \pi_{A_n}(\mathbf{c}) \in \{0,1\}^{d-1}1 \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n-1]\Big\}.$$

Finally, for $i \in \{2, \ldots, n\}$:

$$\mathcal{B}(R_i) := \Big\{\mathbf{c} \mid \exists b \in \{0,1\},\ \bar{b} = 1-b \Big[\pi_{A_i}(\mathbf{c}) \in \{0,1\}^{d-1}b \text{ and } \pi_{A_{i-1}}(\mathbf{c}) \in \{0,1\}^{d-1}\bar{b} \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n] - \{i-1, i\}\Big]\Big\}.$$

There are two possible GYO elimination orders: $(A_1, A_2, \ldots, A_n)$ and $(A_n, A_{n-1}, \ldots, A_1)$. Without loss of generality, let's consider the first.

We make the following claims.

Claim 4. $C := \bigcup_{i \in [n+1]} \mathcal{B}(R_i)$ is a minimal set of boxes that covers $\langle \lambda, \ldots, \lambda\rangle$.

Proof. C covers $\langle\lambda,\ldots,\lambda\rangle$ because $R_2,\ldots,R_n$ constrain $A_1,\ldots,A_n$ to have the same parity while R_1 and R_{n+1} constrain A_1 and A_n to have opposite parities. Moreover, C is a *minimal* set of boxes that covers $\langle\lambda,\ldots,\lambda\rangle$. This is because if any box of $\bigcup_{i\in\{2,\ldots,n\}}\mathcal{B}(R_i)$ is dropped out of C, then A_1 and A_n are no longer constrained to have the same parity. Alternatively, if any box of $\mathcal{B}(R_1)$ or $\mathcal{B}(R_{n+1})$ is dropped, then A_1 and A_n are no longer constrained to have opposite parities. □

Claim 5. In every Tree Ordered Geometric Resolution proof of $\langle\lambda,\ldots,\lambda\rangle$ whose leaves are from C, every box of $\mathcal{B}(R_{n+1})$ must appear $\Omega(N^{(n-1)/2})$ times.

Combined with the fact that $\left|\mathcal{B}(R_{n+1})\right| = \Theta(\sqrt{N})$ (we pick $d = \Theta(\log c)$), proving Claim 5 above proves the theorem.

To complete the proof, we prove Claim 5 by induction on n. The base case is when $n = 2$. By performing resolutions on A_2, we infer the following set of boxes whose A_2 components are λ's:

$$\mathcal{D} := \Big\{\langle x1,\lambda\rangle \quad | \quad x\in\{0,1\}^{d-1}\Big\}.$$

Note that $|\mathcal{D}| = \Theta(\sqrt{N})$. Each box of $\mathcal{B}(R_3)$ is necessary to infer every box of $\mathcal{D}$. (Further, note that to generate $\langle\lambda,\lambda\rangle$ from $\mathcal{B}(R_1)$ we precisely need the boxes in $\mathcal{D}$.)

For the inductive step, we will prove Claim 5 for n assuming it holds for $n-1$. By performing resolutions on A_n, we can infer the following set of boxes whose A_n components are λ's:

$$\mathcal{D} := \Big\{\mathbf{c} \quad | \quad \pi_{A_{n-1}}(\mathbf{c})\in\{0,1\}^{d-1}1 \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k\in[n]-\{n-1\}\Big\}.$$

Note that $|\mathcal{D}| = \Theta(\sqrt{N})$. Because of the resolution order, a box cannot be resolved on any attribute $A_1,\ldots,A_{n-1}$ unless its A_n component is λ. Let $C' := \bigcup_{i\in[n-1]}\mathcal{B}(R_i)\bigcup\mathcal{D}$ be the set of given/inferred boxes whose A_n components are λ's. Because Claim 4 holds for $n-1$, C' is a minimal set of boxes that covers $\langle\lambda,\ldots,\lambda\rangle$. Because Claim 5 holds for $n-1$, every box of $\mathcal{D}$ must appear $\Omega(N^{(n-2)/2})$ times in any proof of $\langle\lambda,\ldots,\lambda\rangle$ from C'. Every box of $\mathcal{B}(R_{n+1})$ must appear in the proof of every box of $\mathcal{D}$. □

G.2 A useful sequence of strings

Before we move ahead we first define a useful collection of sequences of strings.

Definition G.2 ($B_i^{(n)}$). Given an integer $n\geqslant 1$, let $B^{(n)}$ be a sequence of n binary strings $B_1^{(n)},\ldots,B_n^{(n)}$ that are defined as follows:

$$\text{for } i\in[n-1]: \quad B_i^{(n)} = 1^{i-1}0,$$

and

$$B_n^{(n)} = 1^{n-1}.$$

Based on the above definition, $B^{(1)} = (\lambda)$, $B^{(2)} = (0,1)$, $B^{(3)} = (0,10,11)$, $B^{(4)} = (0,10,110,111)$, etc. In other words, $B^{(n)}$ is a binary representation of n values.

We will be interested in these sequences since they have the following nice property:

Lemma G.3. *For any $n\geqslant 1$, $B^{(n)}$ partitions the space $\{0,1\}^n$. Further, there is a unique sequence of resolutions amongst strings/intervals in $B^{(n)}$ that leads to λ.*

Proof. The claim on $B^{(n)}$ being a partition just follows from its definition (as well as our string encoding of dyadic intervals on $\{0,1\}^n$).

For the claim on the sequence note that for every $i\in[n]$, define $R_i^{(n)} = 1^{i-1}$. Note that $B_n^{(n)}(= R_n^{(n)})$ and $B_{n-1}^{(n)}$ resolve to get $R_{n-1}^{(n)}$. More generally, $R_i^{(n)}$ and $B_{i-1}^{(n)}$ resolve to obtain $R_{i-1}^{(n)}$. Note that this sequence results in λ since $R_1^{(n)} = \lambda$. Further, it can be checked that no $B_i^{(n)}$ for $i\in[n-2]$ can be resolved with $B_j^{(n)}$ for $j\neq i$. This implies that the sequence above is the unique way to obtain λ (this can be formally proved e.g. by induction on n). □

G.3 The $\Omega(|C|^{n-1} + Z)$ lowerbound for Ordered Geometric Resolution

Next, we show that the upper bound in Theorem E.10 is tight (even if one used an arbitrary Ordered Geometric Resolution algorithm).

Theorem G.4 (There is a class of input instances on which every Ordered Geometric Resolution algorithm runs in time $\Omega(|C_\square|^{n-1} + Z)$.)**.** *For every integer $n > 1$ and for every integer $c > 0$, there exists a set C of n-dimensional $\Theta(c)$ boxes satisfying the following conditions:*

- *For every box $\mathbf{c} \in C$, $\left|\mathsf{support}(\mathbf{c})\right| \leqslant 3$.*
- *C is a minimal set of boxes that covers $\langle\lambda, \ldots, \lambda\rangle$.* [24]
- *For every fixed order of the n dimensions, every* Ordered Geometric Resolution *proof of $\langle\lambda, \ldots, \lambda\rangle$ consists of $\Omega(|C|^{n-1})$ resolutions.*

Hence, there exists a join query Q whose relations have arity $\leqslant 3$ and whose set of gap boxes $\mathcal{B}(Q)$ is C and whose output is empty, such that every Ordered Geometric Resolution *algorithm runs in time $\Omega(|C|^{n-1})$ on Q.*

Proof. Let $d' := \lceil \log_2 c \rceil$. For each $i \in [n]$, define

$$C_i := \Big\{\mathbf{c} \quad | \quad \exists j \in [n] - \{i\} \left[\pi_{A_i}(\mathbf{c}) = B_j^{(n-1)} \text{ and } \pi_{A_j}(\mathbf{c}) \in \{0,1\}^{d'} \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n] - \{i, j\}\right]\Big\}$$

$$C_i' := \Big\{\langle B_i^{(n)} x_1, x_2, x_3, \ldots, x_n\rangle \quad | \quad \langle x_1, x_2, \ldots, x_n\rangle \in C_i\Big\}$$

$$C := \bigcup_{i \in [n]} C_i'.$$

[25] By definition, $|C| = \Theta(c)$.

We prove the following claim.

Claim 6. For each $i \in [n]$, C_i is a minimal set of boxes that covers $\langle\lambda, \ldots, \lambda\rangle$. Moreover, for every attribute order that ends with A_i, every Ordered Geometric Resolution proof of $\langle\lambda, \ldots, \lambda\rangle$ from C_i consists of $\Omega(|C_i|^{n-1})$ resolutions.

Proof. From Lemma G.3, C_i partitions the universal box $\langle\lambda, \ldots, \lambda\rangle$, which makes C_i a minimal set of boxes that covers $\langle\lambda, \ldots, \lambda\rangle$.

Starting from C_i, by performing all possible resolutions on A_i, we can infer the following set of boxes whose A_i components are λ's

$$\mathcal{D}_i := \Big\{\mathbf{c} \quad | \quad \pi_{A_i}(\mathbf{c}) = \lambda \text{ and } \pi_{A_j}(\mathbf{c}) \in \{0,1\}^{d'} \text{ for all } j \in [n] - \{i\}\Big\}.$$

[26] $\mathcal{D}_i$ is a minimal set of boxes that covers $\langle\lambda, \ldots, \lambda\rangle$. This is because every point in the space is covered by exactly one box of $\mathcal{D}_i$. $|\mathcal{D}_i| = \Omega(|C_i|^{n-1})$, which proves the second part of the claim. □

By Claim 6, for each $i \in [n]$, C_i' is a minimal set of boxes that covers $\langle B_i^{(n)}, \lambda, \ldots, \lambda\rangle$ (and covers nothing outside $\langle B_i^{(n)}, \lambda, \ldots, \lambda\rangle$). Hence, by Lemma G.3, C is a minimal set of boxes that covers $\langle\lambda, \ldots, \lambda\rangle$. Let A_l be the last attribute in the arbitrarily-chosen attribute order. Let's consider the case when $l \neq 1$. (The case of $l = 1$ is going to be very similar.) For any $i \neq j \in [n]$, no box from C_i' can be resolved on A_l with any box from C_j'. Performing resolutions on A_l, we will infer the following boxes from C_l'

$$\mathcal{D}_l' := \Big\{\langle B_l^{(n)} x_1, x_2, x_3, \ldots, x_n\rangle \quad | \quad \langle x_1, x_2, \ldots, x_n\rangle \in \mathcal{D}_l\Big\}.$$

Following the proof of Claim 6, $\mathcal{D}_l'$ is a minimal set of boxes that covers $\langle B_l^{(n)}, \lambda, \ldots, \lambda\rangle$. No box from $C - C_l'$ overlaps with $\langle B_l^{(n)}, \lambda, \ldots, \lambda\rangle$. Hence, each one of the $\Omega(|C|^{n-1})$ boxes of $\mathcal{D}_l'$ is necessary to cover $\langle\lambda, \ldots, \lambda\rangle$. □

[24] This means that (i) the union of all boxes in C covers $\langle\lambda, \ldots, \lambda\rangle$, and (ii) for any box $\mathbf{c} \in C$, $C - \{\mathbf{c}\}$ does not cover $\langle\lambda, \ldots, \lambda\rangle$.

[25] Above, $B_i^{(n)} x_1$ is a *single* string which is the concatenation of $B_i^{(n)}$ and x_1.

[26] Notice that the boxes of $\mathcal{D}_i$ are still in the same n-dimensional space as those of C_i. The A_i components being λ's does not make the boxes of $\mathcal{D}_i$ $(n-1)$-dimensional; it just makes them span the entire A_i dimension.

G.4 The $\Omega(|C|^{w+1} + Z)$ lowerbound for Ordered Geometric Resolution

We begin with a technical lemma.

Lemma G.5. *Given integers $n > 1$ and $d' \geqslant 0$, let C be a set that consists of every n-dimensional box whose support has size 2 and whose two non-λ components share a common suffix after the first d' bits, and this common suffix belongs to $B^{(n-1)}$:*

$$C := \Big\{\mathbf{c} \quad | \quad \exists i \neq j \in [n]\ \exists b \in B^{(n-1)}\Big[\pi_{A_i}(\mathbf{c}), \pi_{A_j}(\mathbf{c}) \in \{0,1\}^{d'} b \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n] - \{i,j\}\Big]\Big\}.$$

Then,

- *C is a minimal set of boxes that covers $\langle \lambda, \ldots, \lambda \rangle$.*
- *Every* Ordered Geometric Resolution *proof of $\langle \lambda, \ldots, \lambda \rangle$ from C consists of $\Omega(|C|^{n/2})$ resolutions.*

Proof. First, we prove that C is a minimal set of boxes that covers $\langle \lambda, \ldots, \lambda \rangle$. Consider an arbitrary but fixed box $\mathbf{b} = \langle x_1, \ldots, x_n \rangle$ where $x_i \in \{0,1\}^{d'}$ for all $i \in [n]$. Consider the following set of boxes that forms a partition of $\mathbf{b}$ (due to Lemma G.3):

$$\mathcal{P}_{\mathbf{b}} = \Big\{\langle x_1 b_1, \ldots, x_n b_n \rangle \quad | \quad b_i \in B^{(n-1)} \text{ for all } i \in [n]\Big\}.$$

Let $C_{\mathbf{b}}$ denote the subset of C whose boxes overlap with $\mathbf{b}$:

$$C_{\mathbf{b}} := \Big\{\mathbf{c} \quad | \quad \exists i \neq j \in [n]\ \exists b \in B^{(n-1)}\Big[\pi_{A_i}(\mathbf{c}) = x_i b \text{ and } \pi_{A_j}(\mathbf{c}) = x_j b \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n] - \{i,j\}\Big]\Big\}.$$

We have n variables $b_1, \ldots, b_n$ (in the definition of $\mathcal{P}_{\mathbf{b}}$). Each variable can take $n-1$ possible values $B^{(n-1)}$. Each box in $C_{\mathbf{b}}$ corresponds to a constraint preventing two of the variables to have the same value. (In particular, a tuple is in a box of $C_{\mathbf{b}}$ if and only if it satisfies the corresponding constraint.) In every assignment of n variables taking $n-1$ values, at least two variables must have the same value. Hence, by the definition of the constraint corresponding to each box in $C_{\mathbf{b}}$, the union of all boxes in $C_{\mathbf{b}}$ covers $\mathbf{b}$. Moreover, if any two of the variables are allowed to have the same value, then there is a feasible assignment of the variables. Hence, for any box $\mathbf{c} \in C_{\mathbf{b}}$, $C_{\mathbf{b}}$ does not cover $\mathbf{b}$ without $\mathbf{c}$.

Now, we prove that $\Omega(|C|^{n/2})$ ordered resolutions are required to infer $\langle \lambda, \ldots, \lambda \rangle$. The n attributes of C are symmetric. WLOG let the attribute order be $(A_1, \ldots, A_n)$. After performing ordered geometric resolutions on A_n eliminating all but the first d' bits, we can infer the following set of boxes:

$$\mathcal{D}_n := \Big\{\mathbf{c} \quad | \quad \pi_{A_n}(\mathbf{c}) \in \{0,1\}^{d'} \text{ and } \Big(\forall i_1 \neq i_2 \in [n-1]\Big)\Big(\exists b_1 \neq b_2 \in B^{(n-1)}\Big)\Big[\pi_{A_{i_1}}(\mathbf{c}) \in \{0,1\}^{d'} b_1 \ \wedge\ \pi_{A_{i_2}}(\mathbf{c}) \in \{0,1\}^{d'} b_2\Big]\Big\}.$$

Notice that $|\mathcal{D}_n| = \Theta(|C|^{n/2})$. Let C_n be the subset of C that contains all boxes whose A_n-components have length $\leqslant d'$:

$$C_n := \Big\{\mathbf{c} \quad | \quad \exists i \in [n-1]\ \exists b \in B^{(n-1)}\Big[\pi_{A_n}(\mathbf{c}), \pi_{A_i}(\mathbf{c}) \in \{0,1\}^{d'} b \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n-1] - \{i\}\Big]\Big\}.$$

$\mathcal{D}_n \cup (C - C_n)$ is a minimal set of boxes that covers $\langle \lambda, \ldots, \lambda \rangle$. This is because boxes of $C - C_n$ correspond to constraints preventing any two of the first $n-1$ variables to have the same value, while boxes of $\mathcal{D}_n$ correspond to constraints preventing the first $n-1$ variables to have $n-1$ different values. (The n-th variable is not constrained.) The union of both constraint types is a minimal set of constraints that rules out all possible assignments of the first $n-1$ variables. □

Next, we show that Theorem 4.7 is tight.

Theorem G.6 (There is a class of input instances on which every Ordered Geometric Resolution algorithm runs in time $\Omega(|C_\square|^{w+1} + Z)$)**.** *For every integer $w > 1$ and for every integer $c > 0$, there is a graph G whose treewidth is w, and there is a set of boxes C whose size is $\Theta(c)$, such that the following properties hold:*

- *For every box $\mathbf{c} \in C$, $\mathsf{support}(\mathbf{c})$ corresponds to some edge in G.*

- *C is a minimal set of boxes that covers $\langle\lambda,\ldots,\lambda\rangle$.*
- *For every elimination order of G that induces a treewidth of w, every* Ordered Geometric Resolution *proof of $\langle\lambda,\ldots,\lambda\rangle$ consists of $\Omega(|C|^{w+1})$ resolutions.*

Hence, there is a join query Q whose (hyper)graph is G and whose set of gap boxes $\mathcal{B}(Q)$ is C and whose output is empty, such that every Ordered Geometric Resolution *algorithm runs in time $\Omega(|C_\square|^{w+1})$ on Q.*

Proof. The vertices of G are divided into two subsets: *primary* and *secondary*. G has $w+1$ primary vertices $A_1,\ldots,A_{w+1}$. For every $i<j\in[w+1]$ and every $k\in[w]$, G has one secondary vertex $B_{i,j,k}$. The edges of G are also divided into primary and secondary. For every $i<j\in[w+1]$, G has one primary edge $\{A_i,A_j\}$. For every $i<j\in[w+1]$ and every $k\in[w]$, G has two secondary edges $\{A_i,B_{i,j,k}\}$ and $\{A_j,B_{i,j,k}\}$.

Let $d':=\lceil\log_2 c\rceil$. For every $i<j\in[w+1]$ and every $k\in[w]$, define $C_{i,j,k}$ as the union of the following two sets

- The set of every box whose support is $\{A_i,B_{i,j,k}\}$, whose $B_{i,j,k}$ component is 0, and whose A_i component ends with the suffix $B_k^{(w)}$ after the first d' bits.
- The set of every box whose support is $\{A_j,B_{i,j,k}\}$, whose $B_{i,j,k}$ component is 1, and whose A_j component ends with the suffix $B_k^{(w)}$ after the first d' bits.

$$C_{i,j,k}:=\Big\{\mathbf{c} \quad | \quad \pi_{A_i}(\mathbf{c})\in\{0,1\}^{d'}B_k^{(w)} \text{ and } \pi_{B_{i,j,k}}(\mathbf{c})=0 \text{ and } \pi_X(\mathbf{c})=\lambda \text{ for all } X\in\text{vars}(Q)-\{A_i,B_{i,j,k}\}\Big\}$$
$$\bigcup\Big\{\mathbf{c} \quad | \quad \pi_{A_j}(\mathbf{c})\in\{0,1\}^{d'}B_k^{(w)} \text{ and } \pi_{B_{i,j,k}}(\mathbf{c})=1 \text{ and } \pi_X(\mathbf{c})=\lambda \text{ for all } X\in\text{vars}(Q)-\{A_j,B_{i,j,k}\}\Big\}.$$

Define C as

$$C:=\bigcup_{i<j\in[w+1]}\ \bigcup_{k\in[w]} C_{i,j,k}.$$

By definition, $|C|=\Theta(c)$.

By performing resolutions on some secondary attribute $B_{i,j,k}$, we can infer the following set of boxes:

$$\mathcal{D}_{i,j,k}:=\Big\{\mathbf{c} \quad | \quad \pi_{A_i}(\mathbf{c})\in\{0,1\}^{d'}B_k^{(w)} \text{ and } \pi_{A_j}(\mathbf{c})\in\{0,1\}^{d'}B_k^{(w)} \text{ and } \pi_X(\mathbf{c})=\lambda \text{ for all } X\in\text{vars}(Q)-\{A_i,A_j\}\Big\}.$$

Notice that every box of $C_{i,j,k}$ is necessary to infer some box of $\mathcal{D}_{i,j,k}$. In particular, if one box of $C_{i,j,k}$ is missing, then we will no longer be able to infer at least one box of $\mathcal{D}_{i,j,k}$. Define $\mathcal{D}$ as

$$\mathcal{D}:=\bigcup_{i<j\in[w+1]}\ \bigcup_{k\in[w]} \mathcal{D}_{i,j,k}.$$

[27] By Lemma G.5, $\mathcal{D}$ is a minimal set of boxes that covers $\langle\lambda,\ldots,\lambda\rangle$. Every box of C is necessary to infer some box of $\mathcal{D}$, which makes $\mathcal{D}$ another minimal set of boxes that covers $\langle\lambda,\ldots,\lambda\rangle$. (We can verify this using the *completeness* of Ordered Geometric Resolution for any fixed attribute order σ. In particular, let's choose σ such that resolution starts with all secondary attributes. From C, we can infer exactly $\mathcal{D}$, which -in turn- is exactly what we need to infer $\langle\lambda,\ldots,\lambda\rangle$. But if some box of C goes missing, we can no longer infer $\mathcal{D}$ or $\langle\lambda,\ldots,\lambda\rangle$.)

Claim 7. Consider some arbitrary but fixed elimination order of G that induces a treewidth of w, and let A_e be the first primary vertex that is eliminated. Then, before A_e can be eliminated, all of its adjacent secondary vertices must be eliminated first.

Proof. A_e is connected with edges to w primary vertices and $w\times w$ secondary vertices. Since non of the w primary vertices could have been eliminated before A_e, all the secondary vertices must be eliminated before A_e, in order to maintain a width of w. □

[27] The boxes of $\mathcal{D}$ are still in the same domain as those of C even though their secondary components $B_{i,j,k}$ are λ's.

WLOG let A_1 be the first primary vertex to-be-eliminated. Before the elimination of A_1, all eliminated vertices were secondary. Whenever we eliminated a secondary vertex $B_{i,j,k}$, we inferred $\mathcal{D}_{i,j,k}$. However, no box from such $\mathcal{D}_{i,j,k}$ could have been resolved any further with any box. This is because boxes of $\mathcal{D}$ can only be resolved on primary vertices, and A_1 is the first one.

After eliminating all the secondary vertices that are adjacent to A_1 (as dictated by Claim 7), we can infer the following set of boxes

$$\mathcal{D}_1 := \bigcup_{1<j\leqslant w+1} \bigcup_{k\in[w]} \mathcal{D}_{1,j,k}.$$

Equivalently, D_1 can be written as

$$\mathcal{D}_1 = \Big\{\mathbf{c} \quad | \quad \exists j \in \{2,\dots,w+1\}\ \exists b \in B^{(w)}\Big[\pi_{A_1}(\mathbf{c}), \pi_{A_j}(\mathbf{c}) \in \{0,1\}^{d'} b \text{ and } \pi_X(\mathbf{c}) = \lambda \text{ for all } X \in \text{vars}(Q) - \{A_1, A_j\}\Big]\Big\}.$$

By performing resolutions on A_1 eliminating all but the first d' bits, we infer the following set

$$\mathcal{E}_1 := \Big\{\mathbf{c} \,|\, \pi_{A_1}(\mathbf{c}) \in \{0,1\}^{d'} \text{ and } \Big(\forall j_1 \neq j_2 \in \{2,\dots,w+1\}\Big)\Big(\exists b_1 \neq b_2 \in B^{(w)}\Big)\Big[\pi_{A_{j_1}}(\mathbf{c}) \in \{0,1\}^{d'} b_1 \ \wedge\ \pi_{A_{j_2}}(\mathbf{c}) \in \{0,1\}^{d'} b_2\Big]\Big\}.$$

[28] From the proof of Lemma G.5, we know that $\mathcal{E}_1 \cup (\mathcal{D} - \mathcal{D}_1)$ is a minimal set of boxes that covers $\langle\lambda,\dots,\lambda\rangle$. Moreover, for every not-yet-eliminated secondary vertex $B_{i,j,k}$, every box of $C_{i,j,k}$ is necessary to infer some box of $D_{i,j,k} \subseteq \mathcal{D} - \mathcal{D}_1$. Noting that $|\mathcal{E}_1| = \Theta(|C|^{w+1})$ completes the proof. □

G.5 The $\Omega(|C|^{n/2} + Z)$ lowerbound for Geometric Resolution

We first prove two structural lemmas that will be useful in proving our lower bound for Geometric Resolution proofs.

Lemma G.7. *Let $\mathcal{A}$ be a set of dyadic boxes. Let $\mathbf{z}$ be a box that can be inferred from $\mathcal{A}$ using r geometric resolutions. Let $\mathcal{A}'$ be a set of boxes that cover the boxes of $\mathcal{A}$ (i.e., for every box $\mathbf{a} \in \mathcal{A}$, there is a box $\mathbf{a}' \in \mathcal{A}'$ such that $\mathbf{a}'$ covers $\mathbf{a}$). Then, a box $\mathbf{z}'$ that covers $\mathbf{z}$ can be inferred from $\mathcal{A}'$ using $r' \leqslant r$ geometric resolutions.*

Proof. By induction. In the base case when $r = 0$, $\mathbf{z}$ must belong to $\mathcal{A}$. Hence, $\mathcal{A}'$ must have some box that covers $\mathbf{z}$.

For the inductive step, we will assume that the lemma holds for $r-1$ and prove it for r. Let $\mathbf{z}$ be a box that can be inferred from $\mathcal{A}$ using r resolutions. Let one of those resolutions be $\mathbf{b} \leftarrow \mathsf{Resolve}(\mathbf{a}_1, \mathbf{a}_2)$ for some $\mathbf{a}_1, \mathbf{a}_2 \in \mathcal{A}$. As a result, $\mathbf{z}$ can be inferred from $\mathcal{B} := \mathcal{A} \cup \{\mathbf{b}\}$ using $r-1$ resolutions. Let $\mathcal{A}'$ be a set of boxes that cover those of $\mathcal{A}$. Let $\mathbf{a}'_1 \in \mathcal{A}'$ cover $\mathbf{a}_1$, and $\mathbf{a}'_2 \in \mathcal{A}'$ cover $\mathbf{a}_2$. We recognize two case:

- If $\mathbf{a}'_1$ and $\mathbf{a}'_2$ can be resolved, then $\mathbf{b}$ must be covered by $\mathbf{b}' \leftarrow \mathsf{Resolve}(\mathbf{a}'_1, \mathbf{a}'_2)$. The boxes of $\mathcal{B}' := \mathcal{A}' \cup \{\mathbf{b}'\}$ cover those of $\mathcal{B}$. Because the lemma holds for $r-1$, a box $\mathbf{z}'$ that covers $\mathbf{z}$ can be inferred from $\mathcal{B}'$ in $\leqslant r-1$ resolutions. Therefore, $\mathbf{z}'$ can be inferred from $\mathcal{A}'$ in $\leqslant r$ resolutions.
- If $\mathbf{a}'_1$ and $\mathbf{a}'_2$ cannot be resolved, then one of them must cover $\mathbf{b}$. Because the lemma holds for $r-1$, a box $\mathbf{z}'$ that covers $\mathbf{z}$ can be inferred from $\mathcal{A}'$ in $\leqslant r-1$ resolutions.

□

Definition G.8. Given a set of boxes $\mathcal{A}$, let $r(\mathcal{A})$ denote the minimum number of geometric resolutions that is sufficient to infer $\langle\lambda,\dots,\lambda\rangle$ from $\mathcal{A}$. (If the union of all boxes in $\mathcal{A}$ does not cover $\langle\lambda,\dots,\lambda\rangle$, then $r(\mathcal{A}) = \infty$.)

Lemma G.9. *Let $\mathcal{A}$ and $\mathcal{B}$ be two sets of boxes that satisfy the following conditions:*

- *$\langle\lambda,\dots,\lambda\rangle$ does not belong to $\mathcal{B}$.*
- *Every geometric resolution between two boxes from $\mathcal{B}$ results in a box that is contained in some box of $\mathcal{A}$.*
- *No box in $\mathcal{B}$ can be geometrically-resolved with any box in $\mathcal{A}$ or with any box that can be inferred from $\mathcal{A}$ through geometric resolution.*

[28] The boxes of $\mathcal{E}_1$ are also assumed to be in the same domain as those of C and $\mathcal{D}$: Their $B_{i,j,k}$ components are hiding as λ's, but they still exist.

Then, $r(\mathcal{A} \cup \mathcal{B}) = r(\mathcal{A})$.

Proof. By definition, $r(\mathcal{A} \cup \mathcal{B}) \leqslant r(\mathcal{A})$. Next, we prove that $r(\mathcal{A} \cup \mathcal{B}) \geqslant r(\mathcal{A})$ by induction on the value of $r(\mathcal{A} \cup \mathcal{B})$. In the base case when $r(\mathcal{A} \cup \mathcal{B}) = 0$, $\langle \lambda, \ldots, \lambda \rangle$ must belong to $\mathcal{A} \cup \mathcal{B}$. Hence, $\langle \lambda, \ldots, \lambda \rangle$ must belong to $\mathcal{A}$ and $r(\mathcal{A}) = 0$.

For the inductive step, we will assume that the lemma holds when $r(\mathcal{A} \cup \mathcal{B}) = k$ for some integer $k \geqslant 0$, and we will prove that it holds when $r(\mathcal{A} \cup \mathcal{B}) = k + 1$. Let $\mathcal{A}$ and $\mathcal{B}$ be two box sets that satisfy the lemma conditions and $r(\mathcal{A} \cup \mathcal{B}) = k+1$. Consider some resolution proof of $\langle \lambda, \ldots, \lambda \rangle$ whose facts are boxes from $\mathcal{A} \cup \mathcal{B}$ and whose number of resolutions is exactly $r(\mathcal{A} \cup \mathcal{B})$. Let $\mathbf{a} \leftarrow \mathsf{Resolve}(\mathbf{a}_1, \mathbf{a}_2)$ be some resolution in this proof such that both $\mathbf{a}_1$ and $\mathbf{a}_2$ belong to $\mathcal{A} \cup \mathcal{B}$. By definition, $r(\mathcal{A} \cup \mathcal{B}) = r(\mathcal{A} \cup \mathcal{B} \cup \{\mathbf{a}\}) + 1$. If both $\mathbf{a}_1$ and $\mathbf{a}_2$ are from $\mathcal{B}$, then $\mathbf{a}$ must be contained in some box of $\mathcal{A}$. By Lemma G.7, $r(\mathcal{A} \cup \mathcal{B}) \leqslant r(\mathcal{A} \cup \mathcal{B} \cup \{\mathbf{a}\})$, which is a contradiction. Because no box from $\mathcal{B}$ resolves with any box from $\mathcal{A}$, both $\mathbf{a}_1$ and $\mathbf{a}_2$ must belong to $\mathcal{A}$. Let $\mathcal{A}' := \mathcal{A} \cup \{\mathbf{a}\}$.

$$r(\mathcal{A} \cup \mathcal{B}) = r(\mathcal{A}' \cup \mathcal{B}) + 1 \geqslant r(\mathcal{A}') + 1 \geqslant r(\mathcal{A}).$$

The first inequality above holds because $r(\mathcal{A}' \cup \mathcal{B}) = k$ and the lemma is assumed to hold for k. The second inequality holds by definition of r. □

We are now ready to argue that Theorem 4.9 is tight.

Theorem G.10 (There is a class of input instances on which every Geometric Resolution algorithm runs in time $\Omega(|\mathcal{C}_\square|^{n/2} + Z)$.)**.** *For every integer $n > 1$ and for every integer $c > 0$, there exists a set C of n-dimensional $\Theta(c)$ boxes satisfying the following conditions:*

- *For every box $\mathbf{c} \in C$, $\left|\mathsf{support}(\mathbf{c})\right| = 2$.*
- *C is a minimal set of boxes that covers $\langle \lambda, \ldots, \lambda \rangle$.*
- *Every* Geometric Resolution *proof of $\langle \lambda, \ldots, \lambda \rangle$ consists of $\Omega(|C|^{n/2})$ resolutions, no matter whether they are ordered or out-of-order.*

Hence, there exists a join query Q whose hypergraph is a clique and whose set of gap boxes $\mathcal{B}(Q)$ is C and whose output is empty, such that every Geometric Resolution *algorithm takes time $\Omega(|C|^{n/2})$ on Q.*

Proof. Let $d' := \lceil 1/2 \log_2 c \rceil$ and C be defined as described in Lemma G.5:

$$C := \Big\{\mathbf{c} \quad | \quad \exists i \neq j \in [n]\ \exists b \in B^{(n-1)} \Big[\pi_{A_i}(\mathbf{c}), \pi_{A_j}(\mathbf{c}) \in \{0,1\}^{d'} b \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n] - \{i, j\}\Big]\Big\}.$$

Note that $|C| = \Theta(c)$ and C is a minimal set of boxes that covers $\langle \lambda, \ldots, \lambda \rangle$.

Lemma G.5 showed that $\Omega(|C|^{n/2})$ ordered resolutions are needed to infer $\langle \lambda, \ldots, \lambda \rangle$ from C. Now we prove that even if we use out-of-order resolutions, we still need $\Omega(|C|^{n/2})$ many of them. For each $l \in \{2, \ldots, n\}$, we define a box set $\mathcal{D}_l$ to consist of every box whose support has size 2 and whose two non-λ components share a common suffix after the first d' bits. This common suffix belongs to $B^{(n-l+1)} - \left\{1^{n-l}\right\}$:

$$\mathcal{D}_l := \Big\{\mathbf{c} \quad | \quad \exists i \neq j \in [n]\ \exists b \in B^{(n-l+1)} - \left\{1^{n-l}\right\} \Big[\pi_{A_i}(\mathbf{c}), \pi_{A_j}(\mathbf{c}) \in \{0,1\}^{d'} b \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n] - \{i, j\}\Big]\Big\}$$

Moreover, for each $l \in \{2, \ldots, n\}$, we define a box set $\mathcal{E}_l$ to consist of every box whose support has size l and whose l non-λ components share a common suffix after the first d' bits. This common suffix is 1^{n-l}:

$$\mathcal{E}_l := \Big\{\mathbf{c} \quad | \quad \exists I \subseteq [n], |I| = l \Big[\pi_{A_i}(\mathbf{c}) \in \{0,1\}^{d'} 1^{n-l} \text{ for all } i \in I \text{ and } \pi_{A_i}(\mathbf{c}) = \lambda \text{ otherwise}\Big]\Big\}$$

From the above definitions, $C = \mathcal{D}_2 \cup \mathcal{E}_2$.

Now, we make the following claim.

Claim 8. For each $l \in \{2, \ldots, n-1\}$, $r(\mathcal{D}_l \cup \mathcal{E}_l) \geqslant r(\mathcal{D}_{l+1} \cup \mathcal{E}_{l+1})$.

Note that $\mathcal{D}_n = \varnothing$ (because $1^0 = \lambda$) and $\mathcal{E}_n = \Big\{\langle x_1, \ldots, x_n\rangle \quad | \quad x_i \in \{0,1\}^{d'} \text{ for all } i \in [n]\Big\}$. Every box in $\mathcal{E}_n$ is necessary to cover $\langle \lambda, \ldots, \lambda\rangle$ (by only using boxes in $\mathcal{E}_n$) and $|\mathcal{E}_n| = \Theta(|C|^{n/2})$. Hence, $r(\mathcal{E}_n) = \Omega(|C|^{n/2})$. Assuming Claim 8 is correct, we have

$$r(C) = r(\mathcal{D}_2 \cup \mathcal{E}_2) \geqslant r(\mathcal{D}_n \cup \mathcal{E}_n) = \Omega(|C|^{n/2}),$$

as desired.

To finish the proof, we now prove Claim 8. The proof of Lemma G.3 implies the following: In every resolution that occurs between two boxes from $\mathcal{D}_l \cup \mathcal{E}_l$, one of the two resolved boxes must belong to $\mathcal{E}_l$ while the other must belong to the following subset of $\mathcal{D}_l$:

$$\mathcal{D}'_l := \Big\{\mathbf{c} \quad | \quad \exists i \neq j \in [n]\Big[\pi_{A_i}(\mathbf{c}), \pi_{A_j}(\mathbf{c}) \in \{0,1\}^{d'} 1^{n-l-1} 0 \text{ and } \pi_{A_k}(\mathbf{c}) = \lambda \text{ for all } k \in [n] - \{i, j\}\Big]\Big\}.$$

This implies that the geometric resolution result must be contained in some box that belongs to $\mathcal{E}_{l+1}$. From the above definition, we have $\mathcal{D}'_l = \mathcal{D}_l - \mathcal{D}_{l+1}$. Hence,

$$r(\mathcal{D}_l \cup \mathcal{E}_l) = r(\mathcal{D}_{l+1} \cup \mathcal{D}'_l \cup \mathcal{E}_l) \geqslant r(\mathcal{D}_{l+1} \cup \mathcal{D}'_l \cup \mathcal{E}_l \cup \mathcal{E}_{l+1}) = r(\mathcal{D}_{l+1} \cup \mathcal{E}_{l+1}).$$

The inequality above holds by definition of r. The second equality holds because of Lemma G.9. □

Corollary G.11. *Let $\mathcal{B}$ be the input set of boxes for Klee's measure problem, n the number of dimensions, and $C_\square$ any optimal certificate for $\mathcal{B}$. Klee's measure problem over the Boolean semiring has a lower bound of $\Omega(|C_\square|^{n/2})$ that holds for all algorithms that are based on* Geometric Resolution.

G.6 The $\omega(|C_\square|^{4/3-\epsilon} + Z)$ lowerbound for arity $\geqslant 3$

The following proposition aims to show that general dyadic boxes are a bit too expressive. In particular, this shows that the restriction of the maximum arity being two in Theorem 4.5 is necessary. (We make use of a well-known complexity theoretic assumption about the hardness of the 3SUM problem [58].)

Proposition G.12 (Hardness of acyclic queries with arity-$\geqslant 3$ relations)**.** *For any $k > 2$, there is a class of (α, β, γ, and Berge)-*acyclic *queries Q whose maximum relation arity is k satisfying the following. Unless the* 3SUM *problem can be solved in sub-quadratic time, there does not exist an algorithm that runs in time $O(|C_\square|^{4/3-\epsilon} + Z)$ for any $\epsilon > 0$ on all input instances. And, there is a class of acyclic queries Q with maximum arity k satisfying the following. Unless the exponential time hypothesis is wrong, no algorithm runs in time $O(|C_\square|^{o(k)})$ on all instances.*

Proof. We first prove the proposition with $k = 3$. We use the 3SUM hardness of the triangle query [58]. Given an instance of the triangle query $R(A,B), S(B,C), T(A,C)$, we create three new three-dimensional relations $W_R(A,B,C)$, $W_S(A,B,C)$ and $W_T(A,B,C)$. The gap boxes for W_R, W_S, W_R are the sets of all gap boxes from R, S, and T, respectively, with appropriate λ's filled in the missing coordinates. For example, if we have a gap box $\langle a, b\rangle$ from R, we insert $\langle a, b, \lambda\rangle$ to $\mathcal{B}(W_R)$. Then, the total size of $\mathcal{B}(W_R), \mathcal{B}(W_S), \mathcal{B}(W_T)$ is linear in the total sizes of R, S, and T, which is linear in the number of tuples from R, S, and T.

Finally, we let Q be any query containing as a sub-query the join $W_R(A,B,C) \bowtie W_S(A,B,C) \bowtie W_T(A,B,C)$, while the rest of the relations from Q can be of arbitrary form. Note that Q can be α, β, γ, or even Berge-acyclic. In the hard instance, we let them contain all possible tuples in their respective domains. Thus, for the rest of the relations in Q, they do not contribute any gap box to the problem. The output of this join is exactly the join of $R \bowtie S \bowtie T$, which cannot be computed in $O(|C_\square|^{4/3-\epsilon} + Z)$-time, following [58]. For $k > 3$, we can pad more λ's to the above reduction.

Using a reduction from unique-k-clique [16] similar to the reduction above, we can also show that no algorithm can run in time $\tilde{O}(|C_\square|^{o(k)})$ for all input instances. The result from [16] states that unique-k-clique does not admit an $O(n^{o(k)})$ algorithm unless the exponential time hypothesis is wrong. □

H Tetris-Reloaded in the GAO-consistent certificate world

When we resolve two boxes $\mathbf{w}_1$ and $\mathbf{w}_2$, the result $\mathbf{w}$ "inherits" some of its components from $\mathbf{w}_1$ and some from $\mathbf{w}_2$ (and some from both). In general, $\mathbf{w}$ is going to be a "mixture" of $\mathbf{w}_1$ and $\mathbf{w}_2$. It might also happen that $\mathbf{w}$ inherits all of its components from $\mathbf{w}_1$ only (or $\mathbf{w}_2$ only). Given a set of M boxes, if we know that no matter how we resolve them, we can never "mix up" more than c of them in a single box, then we cannot perform more than $\tilde{O}(M^c)$ resolutions. The integral cover support lemma (Lemma E.4) was a formalization of this intuitive idea.

Given a hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ and a GAO σ, we can define a measure of how "badly" ordered-geometric-resolution can "mix up" any set of boxes, assuming that those boxes are σ-consistent and their supports correspond to hyperedges in $\mathcal{H}$. We will be referring to this measure as the cover-width of $\mathcal{H}$ that is induced by σ, denoted by $\mathsf{cw}(\sigma)$. Based on this intuition, it is not surprising that Tetris-Reloaded runs in time $\tilde{O}(|C_\square^{\text{gao}}|^{\mathsf{cw}(\sigma)} + Z)$, which is what we are going to prove in this section. This result generalizes earlier ones found in [49].

We start by defining the cover-width of a hypergraph and stating some of its properties.

H.1 Cover-width of hypergraphs

Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be a hypergraph. Let $\sigma = (v_1, \ldots, v_n)$ be an ordering of vertices of this hypergraph. In the context of this paper, this ordering is called a GAO. In traditional graph theory, database, and graphical model applications, an ordering of vertices of a hypergraph is often called an "elimination order." We use GAO instead of elimination order because sometimes we will need our GAO to be a particular elimination order or the reverse of an elimination order in the traditional sense.

For any hypergraph G, let $V(G)$ denote its vertex set, and $E(G)$ denote its edge set. The *union of two hypergraphs* G_1 and G_2, denoted by $G_1 \cup G_2$, is a hypergraph $G = \Big(V(G_1) \cup V(G_2), E(G_1) \cup E(G_2)\Big)$. If G is a hypergraph and v is one of its vertices, then $G - v$ denotes the hypergraph obtained from G by removing v from $V(G)$ and from every edge in $E(G)$. If an empty edge results, we remove the empty edge too.

We define collections $\mathcal{G}_k$, $k \in [n]$, of hypergraphs with respect to σ, recursively, as follows. A hypergraph is identified using the collection of hyperedges that it has; hence, a "system of sets" and a hypergraph are used interchangeably whichever makes more intuitive sense.

To facilitate the recursive definition, define

$$\mathcal{G} = \{\{F\} \mid F \in \mathcal{E}\}.$$

(Here, $\mathcal{E}$ is the set of hyperedges of the original hypergraph $\mathcal{H} = (\mathcal{V}, \mathcal{E})$, the hypergraph of the input query. In words, initially the collection $\mathcal{G}$ consists of $|\mathcal{E}|$ single-edge hypergraphs, one for each member of $\mathcal{E}$, the edge set of $\mathcal{H}$.

For each k from n down to 1, we construct the hypergraph collection $\mathcal{G}_k$ in two steps, then perform two extra steps to prepare for the construction of $\mathcal{G}_{k-1}$.

- **Step 1, initialization:** set $\mathcal{G}_k$ to be

$$\mathcal{G}_k = \{G \mid G \in \mathcal{G} \text{ and } v_k \in V(G)\}.$$

- **Step 2, taking closure:** as long as there are two members $G_1, G_2 \in \mathcal{G}_k$ such that $G_1 \cup G_2 \notin \mathcal{G}_k$, add the hypergraph $G_1 \cup G_2$ to $\mathcal{G}_k$. The construction of $\mathcal{G}_k$ is completed until we can no longer find such two members.

- **Step 3, vertex elimination:** set $\mathcal{G} = \mathcal{G} \cup \mathcal{G}_k$, then for every $G \in \mathcal{G}$, replace G by $G - v_k$.

The *minimum integral edge cover number* of a hypergraph G, denoted by $\rho(G)$, is the minimum number of hyperedges in $E(G)$ that can be used to cover $V(G)$.

The *induced cover-width* of the GAO $\sigma = (v_1, \ldots, v_n)$ with respect to the original hypergraph $\mathcal{H}$, abbreviated by $\mathsf{cw}(\sigma)$, is defined to be

$$\mathsf{cw}(\sigma) := \max_{k \in [n]} \max_{G \in \mathcal{G}_k} \rho(G).$$

The *cover-width* of a hypergraph $\mathcal{H}$ is the minimum induced cover-width over all possible GAOs:

$$\mathsf{cw}(\mathcal{H}) = \min_{\text{GAO } \sigma} \mathsf{cw}(\sigma).$$

The following Propositions show that the notion of cover-width is a natural "interpolation" width notion between β-acyclicity and treewidth. Furthermore, there are classes of hypergraphs for which the cover-width is arbitrarily smaller than the treewidth.

Proposition H.1. *The cover-width of a β-acyclic hypergraph $\mathcal{H}$ is* 1. *Furthermore, a GAO with cover-width* 1 *for $\mathcal{H}$ can be computed in time polynomial in the query complexity.*

Proof. Let $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ be the hypergraph. We showed in [49] that there exists a GAO $\sigma = (v_1, \ldots, v_n)$ which is a nested elimination order. What that means is, for every $k \in [n]$, if we define the following set system:

$$\mathcal{F}_k = \{F \cap \{v_1, \ldots, v_k\} \mid F \in \mathcal{E} \text{ and } v_k \in F\}$$

then the sets in $\mathcal{F}_k$ form a chain, i.e. series of sets where one is contained in the next.

It is sufficient to show that if σ is a nested elimination order, then $\rho(\mathcal{G}_k) = 1$ for every $k \in [n]$, where the $\mathcal{G}_k$ are defined above. It is not hard to see that $\mathcal{G}_k \subseteq 2^{\mathcal{F}_k}$, the power set of $\mathcal{F}_k$. Hence, for every hypergraph $G \in \mathcal{G}_k$, the hyperedges of G form a chain. To cover G, we can simply take the bottom hyperedge on this chain, i.e. $\rho(G) = 1$ for every $G \in \mathcal{G}_k$.

Verifying whether a hypergraph is β-acyclic can be done easily in polytime by an elimination procedure [49], which also yields a nested elimination order if one exists. □

Proposition H.2 ($\mathsf{cw}(\mathcal{H}) \leqslant \mathsf{tw}(\mathcal{H}) + 1$)**.** *If $\mathcal{H}$ is a hypergraph with treewidth w, then its cover-width is at most $w + 1$.*

Proof. Since $\mathsf{tw}(\mathcal{H}) = w$, there exists a GAO $\sigma = (v_1, \ldots, v_n)$ for which the induced width of σ is w. In particular, for every hypergraph $G \in \mathcal{G}_k$, we know $V(G) \subseteq \mathsf{support}(v_k)$, and hence $|V(G)| \leqslant w+1$. (See relation (10).) Any minimal integral edge cover of a hypergraph with at most $w + 1$ vertices must have size at most $w + 1$. Thus, $\rho(G) \leqslant w + 1$ for all $G \in \mathcal{G}_k$ and for all $k \in [n]$. Consequently, $\mathsf{cw}(\mathcal{H}) \leqslant w + 1$ as desired. □

Example H.3 ($\mathsf{cw} \ll \mathsf{tw}$)**.** There are classes of hypergraphs whose cover-width is any given positive integer and whose treewidth is unbounded.

Let c be any positive integer and $\mathcal{H}$ be any hypergraph with $\mathsf{tw}(\mathcal{H}) = c$. Create a new hypergraph $\mathcal{H}'$ from $\mathcal{H}$ by cloning every vertex of $\mathcal{H}$ t times. If an edge of $\mathcal{H}$ contains a vertex v, then the edge will contain t copies of v in $\mathcal{H}'$. It is not hard to see that $\mathsf{cw}(\mathcal{H}') = \mathsf{cw}(\mathcal{H})$ but $\mathsf{tw}(\mathcal{H}')$ is now unbounded as t is arbitrary.

Example H.4 (cw of a Loomis-Whitney query)**.** The LW(n) query is the query whose hypergraph is $\mathcal{H} = \left([n], \binom{[n]}{n-1}\right)$. Note that LW(3) is the triangle query. It is not hard to see that the cover-width of this hypergraph is 2, independent of the GAO. This is (very roughly) because, every hyperedge is missing a unique vertex; and thus any of two of them cover their union.

Example H.5 (cw of a cycle)**.** Let C_n be the cycle of length n. Then, $\mathsf{cw}(C_n) \geqslant 2$ because no matter which GAO we choose, there is a hypergraph in $\mathcal{G}_n$ that has an integral cover number 2. Now, if the cycle's edges are $\{(1,2),(2,3),\cdots,(n-1,n),(n,1)\}$, and let $\sigma = (1,2,\ldots,n)$, then $\mathcal{G}_n$ consists of the following hypergraphs (recall that we identify a hypergraph by its set of hyperedges, so we will not explicitly write down the vertex set of a hypergraph for brevity):

$$\begin{aligned} G_1 &= \{(1,n)\}, \\ G_2 &= \{(n-1,n)\}, \\ G_3 &= \{(1,n),(n-1,n)\} \text{ (this is } G_1 \cup G_2) \end{aligned}$$

As reasoned above, $\rho(G_3) = 2$. Now, after eliminating vertex n, $\mathcal{G}_{n-1}$ consists of the following hypergraphs

$$\begin{aligned}
G'_1 &= \{(n-2, n-1)\} \\
G'_2 &= \{(n-1)\} \text{ (this is } G_2 - n) \\
G'_3 &= \{(1), (n-1)\} \text{ (this is } G_3 - n) \\
G'_4 &= \{(n-1), (n-2, n-1)\} \text{ (this is } G'_1 \cup G'_2) \\
G'_5 &= \{(1), (n-1), (n-2, n-1)\} \text{ (this is } G'_1 \cup G'_3)
\end{aligned}$$

We can see that $\rho(G'_i) \leqslant 2$. Continuing with this process, we can prove by induction that $\mathsf{cw}(C_n) = 2$.

Example H.6 (cw of a clique)**.** The cover-width of a k-clique is $k-1$. We omit the proof.

H.2 GAO-consistent certificate results

The expressiveness of the input gap boxes certainly help significantly reduce the certificate size (see Examples B.7 and B.8). However, as we have seen from the negative result of Proposition G.12, there is a great tension between the expressiveness of the input gap boxes and the runtime of *any* join algorithm.

Example H.7. Consider a query with four attributes, where the GAO is (A_1, A_2, A_3, A_4), and the input relations are $R(A_1, A_2, A_3, A_4)$ and $S(A_2, A_4)$. For two gaps ($\mathbf{w}_1$ from R and $\mathbf{w}_2$ from S) to be resolvable on the last attribute, they must have the form:

$$\begin{aligned}
\mathbf{w}_1 &= \langle t_1 \ , \ t_2 \ , \ t_3 \ , \ x_4 0 \rangle \\
\mathbf{w}_2 &= \langle \lambda \ , \ t_2 \ , \ \lambda \ , \ x_4 1 \rangle
\end{aligned}$$

where t_1, t_2, t_3 are maximal-length strings (The '0' and '1' could have appeared the other way around). The resolution result $\mathbf{w}$ will be:

$$\mathbf{w} = \langle t_1 \ , \ t_2 \ , \ t_3 \ , \ x_4 \rangle$$

Notice that $\mathbf{w}$ is a prefix box of $\mathbf{w_1}$. (See Definition C.2.) In fact, any ordered geometric resolution between two gaps (or prefixes of gaps) from R and S is going to produce a prefix of one of them.

The above example can be made into a formal result. In this section, we show that if we restrict the input gap boxes to be σ-consistent, where σ is the SAO used by Tetris, then there are classes of queries where we can obtain runtimes that are arbitrarily better than the $\tilde{O}(|C_\square|^{w+1} + Z)$ runtime proved in Section 4.4. Of course, the new runtime will be measured on the weaker notion of certificate: the GAO-consistent certificate $C_\square^{\text{gao}}$. GAO-consistent certificates, while weaker than general box certificates, capture a good class of practical algorithms. For example, the Leapfrog Triejoin algorithm [71] implemented in LogicBlox database engine uses only indices from one GAO.

Theorem H.8. *Let $\mathcal{B}$ be a set of σ-consistent boxes for some fixed GAO σ. Let $\mathcal{H}(\mathcal{B})$ be the supporting hypergraph of $\mathcal{B}$, and $\mathsf{cw}(\sigma)$ be the induced cover-width of σ with respect to $\mathcal{H}(\mathcal{B})$. If* Tetris-Reloaded *uses σ as a SAO, it solves* BCP *on input $\mathcal{B}$ in time $\tilde{O}(|C_\square|^{\mathsf{cw}(\sigma)} + Z)$. (Notice that because $\mathcal{B}$ is σ-consistent, $C_\square$ in here will be σ-consistent as well.)*

Corollary H.9 (Tetris generalizes Minesweeper)**.** *Let Q be a join query, σ be a fixed GAO, and suppose that the set of input gap boxes $\mathcal{B}(Q)$ is σ-consistent. If* Tetris-Reloaded *uses σ as a SAO, it solves Q in time $\tilde{O}(|C_\square^{gao}|^{\mathsf{cw}(\sigma)} + Z)$.*

Proof of Theorem H.8. WLOG, assume $\sigma = (A_1, \ldots, A_n)$. Let $c = \mathsf{cw}(\sigma)$. The proof strategy is as follows. We apply Lemma E.4 by showing that, for every gap box resolvent $\mathbf{w}$, there exists an integral cover $F_1, \ldots, F_c$ for $\mathsf{support}(\mathbf{w})$ such that $\mathbf{w}$ is supported on $\{F_1, \ldots, F_c\}$. To get the desired integral cover, we equip $\mathbf{w}$ with a hypergraph $G = (V = \mathsf{support}(\mathbf{w}), E)$ such that for *every* edge F of this hypergraph, $\mathbf{w}$ is supported on F by some input gap box, and that the minimum integral cover for the hypergraph has size at most c. We obtain such hypergraph by induction. If $\mathbf{w}$ was the result of resolving $\mathbf{w}_1$ with $\mathbf{w}_2$, then the hypergraph G will be constructed from the hypergraphs of $\mathbf{w}_1$ and $\mathbf{w}_2$, inductively.

In the base case, if $\mathbf{w}$ is an input gap box, then the hypergraph for $\mathbf{w}$ is a single-edge hypergraph $G = (V, E)$ with $V = \mathsf{support}(\mathbf{w})$ and $E = \{V\}$. Clearly $\mathbf{w}$ is supported on V.

Now, suppose $\mathbf{w}_1$ and $\mathbf{w}_2$ are two gap boxes or gap box resolvents satisfying the following conditions. There are two hypergraphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ for which $V_i = \mathsf{support}(\mathbf{w}_i), i \in [2]$, and for every $F \in E_i$, $\mathbf{w}_i$ is supported on F by some input gap box.

Let $\mathbf{w} = \mathsf{Resolve}(\mathbf{w}_1, \mathbf{w}_2)$, and suppose the resolution is on attribute A_k for some $k \in [n]$. We consider two cases.

Case 1. The A_k-component of $\mathbf{w}$ is not λ. It is *crucial* that the gap boxes are σ-consistent, so every component of $\mathbf{w}_1$ and $\mathbf{w}_2$ after A_k is λ and every component before A_k is either λ or full-length. (Recall Definition 3.10.) Consequently, for every subset $S \subseteq \mathsf{support}(\mathbf{w}_i)$, we have $\pi_S(\mathbf{w}_i) \subseteq \pi_S(\mathbf{w})$. It follows that the hypergraph $G = G_1 \cup G_2$ satisfies the condition that $\mathsf{support}(\mathbf{w}) = V(G)$ and $\mathbf{w}$ is supported on each hyperedge in $E(G)$.

Case 2. The A_k-component of $\mathbf{w}$ is λ. In this case, we equip $\mathbf{w}$ with the hypergraph $(G_1 \cup G_2) - A_k$.

Finally, by the definition of $\mathsf{cw}(\sigma)$, for each $\mathbf{w}$ and each hypergraph G that $\mathbf{w}$ is equipped with, the minimum integral cover of G has size at most $\mathsf{cw}(\sigma)$. And since $\mathbf{w}$ is supported on any integral cover, the theorem is proved. □

Corollary H.10 (Minesweeper's guarantee)**.** *If Q is β-acyclic, then there exists a SAO σ for which* Tetris-Reloaded *runs in time $\tilde{O}(|C_\square^{gao}| + Z)$, given that the set of input gap boxes $\mathcal{B}(Q)$ is σ-consistent.*

Proof. This follows from the above theorem and Proposition H.1. □

H.3 Cutset cover-width

The notion of cover-width was an attempt to measure how β-acyclic a hypergraph is. A β-acyclic hypergraph has cover-width exactly 1. In probabilistic graphical model (PGM) inference [59] and constraint satisfaction problem (CSP) solving [21], researchers have noticed for a long time that most inference or CSP problems are easy on α-acyclic hypergraphs. (Note that, as was shown in [49], in the certificate world the boundary has moved from α-acyclicity to β-acyclicity.)

In PGM inference or CSP, one way to measure the degree of α-acyclicity of a hypergraph is to count the minimum number of vertices one has to remove to make the hypergraph α-acyclic. This idea gives rise to the *cycle cutset conditioning* algorithms [20, 59]. The runtime is then multiplied by an exponential factor in the number of vertices removed. We do not necessarily need a resulting α-acyclic hypergraph. We might want to stop at a hypergraph of small treewidth, at which point a treewidth-based search algorithm takes over. This is the idea of w-cutset conditioning algorithms [11].

Following the same line, we can define a notion of *cutset c-cover-width*, which is the minimum number x of vertices, say $X \subseteq V$, we have to remove so that the resulting hypergraph $H - X$ has cover-width equal to c. In that case, it is easy to see that Tetris-Reloaded runs in time $\tilde{O}(|C_\square^{\text{gao}}|^{x+c} + Z)$ given the correct SAO. All we have to do is to put the removed attributes in X in front of the SAO, and the best attribute ordering with respect to the cover-width of the residual graph $H - X$ in the end of the SAO. Then, we apply Lemma E.4 as follows. For a given witness $\mathbf{w}$, we construct an integral cover of its support set by having a singleton set for each attribute in $\mathsf{support}(\mathbf{w}) \cap X$, and the usual integral cover on $\mathsf{support}(\mathbf{w}) \cap (V - X)$.

However, it is easy to prove the following, essentially saying that cutset conditioning does not help our cause in this problem setting.

Proposition H.11. *If a hypergraph H has cutset c-cover-width x, then $\mathsf{cw}(H) \leqslant c + x$.*

Remark H.12. What does help, however, is that the dyadic segments in the X components can be general dyadic segments, because in the integral cover we use only singleton sets as the supports. Hence, this idea of cutset cover-width might still be useful when we know there is a small subset of attributes on which the input gap boxes store general dyadic segments.

I From geometry to logic: DNF certificates and connection to DPLL with clause learning

The bit-string encoding of dyadic intervals leads to a very natural idea (in hind sight), and opens up a large number of intriguing questions regarding the nature of database indices and the join operation on them. In this section, we briefly touch upon some of the research directions and the questions that Tetris pointed to.

We discuss an alternative formulation, turning the problem from geometry to logic. We do so by explaining how to store data as DNF-formulas, replacing the gap boxes by the more expressive DNF terms. This encoding/indexing leads to the very natural notion of DNF-certificates, which can be *a lot* smaller than geometric certificates, at a price. Then, we explain how syntactically one can view Tetris as doing a special form of DPLL with clause learning.

This section is partly speculative. In some sense we "close the loop": SAT is a special case of constraint satisfaction problem, which is equivalent to conjunctive query evaluation, which Minesweeper showed to be geometric, which Tetris expanded, which led us back to SAT.

I.1 DNF database indices

Fix an input relation R of arity k. Again for simplicity, assume the domain of all attributes are of size $D = 2^d$. We encode each input tuple $\mathbf{t} = (t_1, \ldots, t_k) \in R$ with a DNF term in a very natural way. There are dk (bit) variables $x_1, \ldots, x_{dk}$, of which d variables are used to represent each t_i. For example, for $k = 2$ and $d = 3$, and

$$\begin{aligned} t_1 &= 2\ (\texttt{010} \text{ in binary}) \\ t_2 &= 6\ (\texttt{110} \text{ in binary}), \end{aligned}$$

we have a truth assignment – with notation overloading –

$$\mathbf{t}(x_1, \ldots, x_6) = (\texttt{010}, \texttt{110}).$$

This is the unique satisfying assignment to the *DNF-term* (or *conjunctive clause*)

$$\bar{x}_1 \wedge x_2 \wedge \bar{x}_3 \wedge x_4 \wedge x_5 \wedge \bar{x}_6.$$

Definition I.1 (Tuple DNF-formula)**.** Any k-ary relation R is simply a *DNF formula* (a disjunction of DNF-terms), where each DNF-term has exactly dk variables. Each term represents a tuple in the relation. The set of tuples in R is precisely the set of truth assignments satisfying the DNF formula. We will refer to this formula as $\varphi(R)$, and call it the *tuple DNF formula* encoding R.

The above encoding has size $\tilde{O}(|R|)$, and one can certainly envision building an index for the relation that way; though it is not quite clear what we gain from doing so. Next, we draw inspiration from Minesweeper and Tetris: we would like the index to be able to return "gaps" representing a region of space where no tuple from R resides.

The first natural idea is to use the complement of the tuple DNF formula to represent the gaps. This complement is a *CNF-formula*, which is a conjunction of clauses. This representation would represent all possible gaps at once. The problem, of course, is that this representation has size $\Omega(kd|R|)$. Consequently, if we use such a CNF-formula to answer a probe, it would be the same as transmitting back the entire relation, defeating the purpose of a probe. Note that in this setting, a probe is simply a truth assignment of kd variables. (The projection of the higher dimensional probe point down to this relation's attributes.)

Recall that the gap boxes cover the non-input-tuples in the union-sense: they can overlap. The gap boxes represent DNF-like formulas. This observation leads to the next natural idea. We can design DNF formulas to represent gaps too. Each term of a DNF formula has at most dk variables, and thus to respond to a probe we can return a few terms that the probe satisfies without the space explosion.

The dyadic gap boxes are one type of gap DNF terms we are looking for. Let us start with a couple of examples.

Example I.2. Consider the following relation $R(A) = \{1, 5, 11\}$ where $\mathbf{D}(A) = \{0, 1, \ldots, 15\}$, which means $d = 4$. The tuple DNF formula is

$$\varphi(R) = (\bar{x}_1 \wedge \bar{x}_2 \wedge \bar{x}_3 \wedge x_4) \vee (\bar{x}_1 \wedge x_2 \wedge \bar{x}_3 \wedge x_4) \vee (x_1 \wedge \bar{x}_2 \wedge x_3 \wedge x_4).$$

The truth assignments not in R form gaps, which are as follows.

- Gap 1: {0000}, represented by the DNF term

$$\bar{x}_1 \wedge \bar{x}_2 \wedge \bar{x}_3 \wedge \bar{x}_4$$

- Gap 2: {0010, 0011, 0100}, represented by the DNF formula

$$(\bar{x}_1 \wedge \bar{x}_2 \wedge x_3) \vee (\bar{x}_1 \wedge x_2 \wedge \bar{x}_3 \wedge \bar{x}_4).$$

- Gap 3: {0110, 0111, 1000, 1001, 1010}, represented by the DNF formula

$$(\bar{x}_1 \wedge x_2 \wedge x_3) \vee (x_1 \wedge \bar{x}_2 \wedge \bar{x}_3) \vee (x_1 \wedge \bar{x}_2 \wedge x_3 \wedge \bar{x}_4).$$

- Gap 4: {1100, 1101, 1110, 1111}, represented by the DNF formula

$$(x_1 \wedge x_2).$$

It should be clear that the DNF-terms above correspond precisely to the set of all dyadic boxes. And it should also be clear that gap DNF-formulas are much more powerful than the dyadic boxes. For example, it is possible to merge the second DNF-term of Gap 2 with the DNF-term from Gap 1 to form a DNF term $(\bar{x}_1 \wedge \bar{x}_3 \wedge \bar{x}_4)$. We will get back to this crucial point later.

Example I.3 (The interleaving case)**.** Consider the hard instance for Minesweeper and Tetris, where we want to compute the join $R(A) \bowtie S(A)$ with $R(A) = \{0, 2, \ldots, 2^k - 2\}$, $S(A) = \{1, 3, \ldots, 2^k - 1\}$. Minesweeper and Tetris work in the same way for this example: they run in time $\Omega(|R| + |S|)$.

However, if by magic the database index infers more about the relations and represents the DNF formulas much more succinctly, then join algorithms can run a lot faster. The complement of $R(A)$ can be represented by the DNF-formula

$$\bar{R} = x_d,$$

and the complement of $S(A)$ is

$$\bar{S} = \bar{x}_d.$$

And the certificate $\bar{x}_d \vee x_d$ certifies that the output is empty.

Proposition I.4. *Every dyadic box (in any index order) can be presented by a DNF-term. The converse does not hold.*

Proof. The forward direction is obvious. The interleaving example above shows that the converse does not hold: the DNF-term x_d cannot be represented by a single dyadic box. □

Definition I.5 (Gap DNF formula)**.** Given a relation R, a gap DNF formula is a DNF formula $\bar{\varphi}(R)$ such that $\mathbf{t} \notin R$ iff $\mathbf{t}$ satisfies $\bar{\varphi}(R)$. A gap DNF formula is *non-redundant* if none of its term logically infers another.

Definition I.6 (DNF index)**.** A DNF index is a data structure storing tuples in a relation R such that the storage maintains a gap DNF formula. Given a queried tuple, the index either returns *YES*, the tuple belongs to R, or *NO* the tuple does not belong to R. In the NO case, a set of gap DNF terms from the gap DNF formula which the tuple satisfies are returned as evidence.

We leave the *many* tradeoffs involved in building such a database index for an (exciting) future work. The above definition is necessarily vague. As far as we know there is no such index in the database literature. (*Bitmap/Variant indices*, widely implemented in database management systems [54, 55, 67], are close to the spirit of DNF indices.)

We would like to emphasize, however, that what Minesweeper and Tetris pointed to is the following: ordered indices such as B-trees, tries, or even hash tables (up to a log-factor loss) can effectively be viewed as DNF indices.

To briefly touch upon the huge space of tradeoffs involved in designing such an index, let us even leave aside the all-important question of how to efficiently maintain and update relations built using a DNF index. The *theoretical*

question of which gap DNF-formula to maintain is already interesting and difficult. There are *many* gap DNF-formula for a given relation R. The gap DNF formula corresponding to the dyadic gap boxes has at least as many terms as $|R|$, one for each gap. As the interleaving example shows, this representation might be wildly redundant. It might make sense to find a DNF-formula with the minimum number of terms so that the storage is compact.

Finding a minimum DNF-formula representing a given relation R is known to be NP-hard [3, 46]. It is easy to see that MIN-DNF is a special case of SET COVER, and thus admits a $\log N$-approximation algorithm. Here, N is the size of the truth table. In other words, we can approximate MIN-DNF to within a factor of dk, where k is the number of attributes of R, and d is the number of bits to represent each attribute value. This bound is almost tight, since there is a $\gamma > 0$ for which $(\log N)^{\gamma}$-approximation is not possible (modulo a well-known complexity theoretic assumption) [3].

I.2 DNF-certificate and DPLL with Clause Learning

In the previous section, we have established that a DNF-index can generalize indices that store dyadic gap boxes; the next natural step is to define the notion of DNF-certificate for a join query. As defined above, each relation can abstractly be viewed as a gap DNF-formula, containing many DNF terms. A tuple does not belong to the relation iff it satisfies some term of the gap DNF-formula. We use $\mathcal{D}(R)$ to denote the set of all DNF terms of the gap DNF formula for relation R.

Definition I.7 (DNF-certificate)**.** A *DNF-certificate* for an instance of the join query Q is a collection C_{dnf} of DNF-terms, where

$$C_{\text{dnf}} \subseteq \bigcup_{R \in \text{atoms}(Q)} \mathcal{D}(R),$$

such that a tuple $\mathbf{t}$ (i.e. a truth assignment) satisfies C_{dnf} iff it is *not* an output tuple of Q.

If the DNF-indices are built such that the number of DNF terms in $\mathcal{D}(R)$ which a probe point satisfies is $\tilde{O}(1)$, then we can discover an optimal C_{dnf} using a Minesweeper-like algorithm in the same way that Tetris-Reloaded was designed. Each time we probe into a relation, it will return either (1) YES, meaning the probe belongs to the relation, or (2) NO, along with a set of gap DNF terms, representing a set of gap DNF terms that the probe satisfies.

The set of all gap DNF terms discovered so far can be stored in a data structure. Then, the next probe point (or negative witness in Tetris' sense) is computed from this data structure. It is a tuple that does not satisfy any of the DNF terms stored in the data structure. In other words, it is a satisfying truth assignment to a CNF formula. (The complement of the gap DNF terms.)

Hence, we are entering the realm of SAT solvers. A SAT-solver takes a CNF-formula and either provides a *refutation proof* that the formula is not satisfiable, or a satisfying assignment. Most known SAT algorithms are based on variations of the Davis-Putnam procedures [18], which is resolution-based, or the Davis-Putnam-Logemann-Loveland (DPLL) algorithm [17], which is pure backtracking search.

The original Davis-Putnam procedure [18] is based on resolution, and suffers from the memory explosion problem. DP-resolution can be cast as a variable-elimination algorithm [24], which has the identical structure as the variable-elimination algorithm for graphical model inference. If the variable ordering is chosen so that the induced treewidth of the sequence is small, then it might be faster then DPLL [24]. DP-resolution can be very bad [30] in theory, compared to other proofs, if we aim to minimize the proof size. The *DPLL-algorithm* [17] is search-based. The pure form of this algorithm eliminates the exponential memory requirement of DP-resolution. DPLL can be viewed as DP with unit resolution. Most practically efficient SAT-solving algorithms use DPLL with more sophisticated forms of *clause-learning* [26, 48, 65, 66, 73] (see [41] for a nice survey). The idea is to insert back into the CNF formula a new clause that the algorithm has learned during the search, in order to "cache" some of the computation performed thus far.

In summary, one can envision applying known SAT solving algorithms to compute joins with runtime proportional to the optimal DNF certificates. General DNF certificates are extremely expressive, and thus the general problem is probably too difficult. As we have mentioned earlier, we leave open the question of how a DNF index can be built and maintained efficiently. We leave this direction of inquiry to a future work.

There are a couple of observations that, in hind sight, relate Tetris to DPLL with clause learning.

Tetris is DPLL with clause learning If we view the input box $\mathbf{b}$ to Tetris as a partial truth assignment, then the partition of $\mathbf{b}$ into $\mathbf{b}_1$ and $\mathbf{b}_2$ is simply the assigned value of the *next* variable. The box $\mathbf{w}_1$ or $\mathbf{w}_2$ can be thought of as encodings of a learned clause or a base clause, where we found that the current partition assignment $\mathbf{b}_1$ or $\mathbf{b}_2$ is in conflict with the CNF clauses. If $\mathbf{w}_1$ or $\mathbf{w}_2$ contains $\mathbf{b}$, i.e. $\mathbf{b}$ violates one of those clauses, then the algorithm backtracks. Otherwise, the partial assignment $\mathbf{b}$ violates the resolved clause $\mathbf{w}$. This resolution is of DP-style. The newly learned clause $\mathbf{w}$ is inserted back into the knowledge base. Furthermore, since the algorithm continues running after a satisfying truth assignment is found, it should properly be thought of as #DPLL.

Worst-case optimal algorithms are pure DPLL with a fixed variable ordering. Another interesting observation is that of Corollary D.3. Tetris does not need to cache resolvents (i.e. it does not need to insert back learned clauses) in order to achieve the worst-case AGM bound for the input query. In hind sight this was also obvious from observing the algorithms from [50] and [71]: those algorithms do not need caching at all.

On the other hand, in order to achieve the fractional hypertree width bound or Yannakakis linear runtime for α-acyclic queries, pure backtracking search is not sufficient. This holds true for the certificate world too, where caching is a must.

J General Resolution can be more powerful than Geometric Resolution

J.1 The example

We now show that the hard examples in the proof of Theorem G.10 can be solved with $\tilde{O}(|C|)$ many general resolutions. Contrast this with the result in Theorem G.10, which states that for the same examples, any geometric resolution scheme needs to make $\Omega(|C|^{n/2})$ resolutions. Thus, this will show that geometric resolution is strictly less powerful than general resolution.

We begin with the case of $n = 3$. For completeness, we restate the hard instance for $n = 3$ from the proof of Theorem G.10 here. We need an instance for the triangle query $R(A,B) \bowtie S(B,C) \bowtie T(A,C)$. Let $M = 2^d$ be an integer for another integer parameter $d \geqslant 1$ and let O_M and $\mathcal{E}_M$ denote the set of odd and even numbers respectively in $\{0,\ldots,M-1\}$. Then the following is a hard instance for geometric resolution:

$$\begin{aligned} R(A,B) &= O_M \times \mathcal{E}_M \cup \mathcal{E}_M \times O_M, \\ S(B,C) &= O_M \times \mathcal{E}_M \cup \mathcal{E}_M \times O_M, \\ T(A,C) &= O_M \times \mathcal{E}_M \cup \mathcal{E}_M \times O_M. \end{aligned}$$

Before we proceed we need to state how we encode the domains so that we can express the gaps as clauses (since we want to use general resolution on them). We will follow the encoding from Section I. In particular, for attributes A, B and C we define d boolean variables each: $a_0,\ldots,a_{d-1}$, $b_0,\ldots,b_{d-1}$ and $c_0,\ldots,c_{d-1}$. Then for every constant $i \in \mathbf{D}(A) = \{0,1\}^d$, we associate a shorthand notation A_i for the conjunctive clause that naturally encodes the binary representation of i. In particular, for $i = \sum_{j=0}^{d-1} i_j \cdot 2^j$ (where $i_j \in \{0,1\}$), the corresponding conjunctive clause will have the literal $\neg a_j$ if $i_j = 0$ and a_j otherwise. Note that $\neg A_i$ is a valid (disjunctive) clause. We similarly define B_i and C_j for $i \in \mathbf{D}(B) = \{0,1\}^d$ and $j \in \mathbf{D}(C) = \{0,1\}^d$. We would like to stress here that we are just using A_i, B_i, C_i as notational macros and there is not necessarily a semantic meaning to this notation. (See Section J.2 for why this might matter in general.)

Before we proceed with the encoding of gaps, we record the following simple observation.

Lemma J.1. *Given the clauses $\neg A_i$ for $i \in O_M$, one can generate the clause $\neg a_0$ with $O(M)$ general resolutions. Similarly, given clauses $\neg A_i$ for $i \in \mathcal{E}_M$, one can generate the clause a_0 with $O(M)$ general resolutions. (Similar results hold for clauses from attributes B and C.)*

Proof. This follows by noting that in general resolution one can essentially ignore the variable a_0 in the clauses and then make $O(M)$ resolutions to generate $\langle \lambda,\ldots,\lambda \rangle$ (on the variables $a_1,\ldots,a_{d-1}$). The proof follows by noting that the values $\neg A_i$ for $i \in O_M$ ($i \in \mathcal{E}_M$ resp.) contain the literal $\neg a_0$ (a_0 resp.). □

We are now ready to state the gaps from the relations above in terms of the notation defined above. Let us consider the gaps from relation R: note that we have a gap rectangle $\langle i, j, \lambda\rangle$ for $(i, j) \in O_M \times O_M \cup \mathcal{E}_M \times \mathcal{E}_M$. In particular, for each such pair (i, j), we have the following conjunctive clause

$$\mathcal{G}_{R,i,j} = A_i \wedge B_j.$$

Note that the negation of the clause is a normal disjunctive clause:

$$\neg\mathcal{G}_{R,i,j} = \neg A_i \vee \neg B_j.$$

Similarly we can define the clauses corresponding to gaps in S (denoted by $\mathcal{G}_{S,i,j}$) and in T (denoted by $\mathcal{G}_{T,i,j}$). Note that checking if the gaps cover $\langle\lambda, \lambda, \lambda\rangle$ is equivalent to checking that the following CNF has no satisfying solutions:

$$\left(\bigwedge_{(i,j)\in O_M\times O_M\cup\mathcal{E}_M\times\mathcal{E}_M} \neg\mathcal{G}_{R,i,j}\right) \wedge \left(\bigwedge_{(i,j)\in O_M\times O_M\cup\mathcal{E}_M\times\mathcal{E}_M} \neg\mathcal{G}_{S,i,j}\right) \wedge \left(\bigwedge_{(i,j)\in O_M\times O_M\cup\mathcal{E}_M\times\mathcal{E}_M} \neg\mathcal{G}_{T,i,j}\right). \tag{22}$$

We will now argue that for the set of clauses above, general resolution is strictly more powerful than geometric resolution.

Lemma J.2. *Starting with* (22) *one can derive a contradiction with $O(M^2)$ general resolutions, while doing the same takes $\Omega(M^3)$ geometric resolutions.*

We thank Paul Beame for telling us the proof below and kindly allowing us to use it here.

Proof. The lower bound follows from the proof of Theorem G.10 so we focus on the upper bound in this proof. Consider the following sequence of (general) resolutions:

1. Fix an $o \in O_M$. Then note that with $O(M)$ resolutions on clauses $\neg A_o \vee \neg B_{o'}$ for $o' \in O_M$, we can generate the clause $\neg A_o \vee \neg b_0$. (This essentially follows from Lemma J.1.) In particular, with $O(M^2)$ general resolutions we can generate the clauses

   $$\neg A_o \vee \neg b_0 \text{ for every } o \in O_M.$$

   Again applying Lemma J.1 to the above set of clauses with further $O(M)$ general resolutions, one can generate the clause

   $$\neg a_0 \vee \neg b_0. \tag{23}$$

   Similarly with $O(M^2)$ general resolutions on the clauses $\neg\mathcal{G}_{R,i,j}$ for $(i, j) \in \mathcal{E}_M \times \mathcal{E}_M$, we get the following clause:

   $$a_0 \vee b_0. \tag{24}$$

2. Using analogous argument as above to the gaps from S and T we can with $O(M^2)$ general resolutions generate the clauses:

   $$\neg b_0 \vee \neg c_0. \tag{25}$$

   $$b_0 \vee c_0. \tag{26}$$

   $$\neg a_0 \vee \neg c_0. \tag{27}$$

   $$a_0 \vee c_0. \tag{28}$$

3. Now with four more resolutions we can generate the clauses a_0 and $\neg a_0$, which will generate the required contradiction. (Indeed to generate a_0, resolve (24) with (25) and resolving the resulting clause $a_0 \vee \neg c_0$ with (28). Similarly, $\neg a_0$ can be generated from the other three clauses.)

The proof is complete by noting that the above steps use $O(M^2)$ general resolutions to generate a contradiction, as required. □

The previous lemma can be generalized to any $n \geqslant 3$. In particular, for every $n \geqslant 3$, the lowerbound example that needs $\Omega(|C|^{n/2})$ geometric resolutions in Theorem G.10 can actually be solved within $\tilde{O}(|C|)$ general resolutions. To see this, recall that a dyadic box is an n-tuple of binary strings of length $\leqslant d$. In the logic framework, each bit corresponds to a Boolean variable and each dyadic box corresponds to a conjunctive clause: variables that appear in this clause are those whose corresponding bits appear in the corresponding dyadic box. (e.g. If some string has length $< d$, then the variable corresponding to the last bit does not appear in the clause.) The negation of a dyadic box is nothing but a disjunctive clause.

To solve the lowerbound example within $\tilde{O}(|C|)$ general resolutions, all we have to do is start with making all resolutions on the Boolean variables corresponding to the first d' bits of each one of the n strings. (Recall the parameters and example from the proof of Theorem G.10.) Within $\tilde{O}(|C|)$ such resolutions, we can infer $\tilde{O}(1)$ clauses that do not contain any one of those $d' \times n$ variables. Within $\tilde{O}(1)$ more resolutions, we can generate a contradiction.

J.2 Encoding Issues

In the previous section, we used the natural binary encoding for encoding each element of the domain as a clause. Now we consider another natural encoding with the goal of pointing out that for our setting the choice of encoding matters. In particular, for every element i of domain $\mathbf{D}(A)$, we have a variable A_i. (Note that now unlike the previous section this is not just a syntactic variable.) One also has to add a clause $\vee_{i \in \mathbf{D}(A)} A_i$ to explicitly define the domain. This essentially corresponds to a unary representation of the domain and is a representation in proof complexity that has been studied before (see e.g. [10]). Below we state the reasons why the binary encoding that we use above is more reasonable than this unary encoding for our purposes.

First, we note that if we assume that all the domains are $\{0,1\}^d$, then since in our setting we have n attributes, any clause can have at most dn literals. This implies that a single general resolution (and obviously geometric resolution) can be implemented in $O(dn)$ time, which is $\tilde{O}(1)$ in our setting. On the other hand, clauses in the unary representation can have $\Theta(nN)$ many literals, which is significantly large in many of our settings. In traditional worst-case proof complexity, this is fine since an extra factor of $O(N)$ in the run time is not a big deal. However, in our beyond worst-case results, this extra factor is too prohibitive.

Second, we note the fact that unary encoding can lead to long clauses, which allows general clauses to encode arbitrary subsets of the domain. By contrast this is not possible in the binary encoding setting (which can only encode polynomial many subsets). In the general attribute case this corresponds to the fact that general resolution can encode combinatorial rectangles/boxes while in the geometric resolution case, we can only encode geometric boxes/resolutions.

Finally, we note that for the case when we have only one attribute (say A and hence we only have the variables $a_0, \ldots, a_{d-1}$), general resolution can be lossy in the following sense. Consider the case where we have gaps at 7 and 11 (i.e. we have $d = 4$). The negations of these gaps can be encoded as

$$a_3 \vee \neg a_2 \vee \neg a_1 \vee \neg a_0 \text{ and } \neg a_3 \vee a_2 \vee \neg a_1 \vee \neg a_0.$$

Note that the resolution of the two clauses above results in $\neg a_1 \vee \neg a_0$, which corresponds to all elements in $\{0,1\}^4$ that do not have their least two significant bits as 0. Note that this set is a *strict* subset of values that are not 5 and not 7. By contrast, when we perform geometric resolution on two clauses corresponding to negation of dyadic intervals, the resulting clause corresponds exactly to the negation of the union of the two dyadic intervals.[29] Thus, in some sense binary encoding is more suited for geometric resolution than to general resolution.

[29] We note however that for two attributes geometric resolution is also similarly lossy.